

Automation and Polarization

Daron Acemoglu

Massachusetts Institute of Technology

Jonas Loebbing

Ludwig Maximilian University of Munich

We develop an assignment model of automation. Each of a continuum of tasks of variable complexity is assigned to either capital or one of a continuum of labor skills. We characterize conditions for interior automation, whereby tasks of intermediate complexity are performed by capital. Interior automation arises when low-skill wages are low and effective cost of capital in low-complexity tasks is high. Minimum wages make interior automation less likely. Higher capital productivity causes employment and wage polarization, changes the skill premium non-monotonically, and reduces the real wage of workers with comparative advantage profiles close to that of capital.

I. Introduction

Automation technologies—including specialized software tools, computerized production equipment, and industrial robots—have been spreading rapidly throughout the industrialized world. For example, the share of information processing equipment and software in overall investment

We thank Juanita Jaramillo for excellent research assistance; David Autor, Chris Busch, Moritz Drechsel-Grau, Pascual Restrepo, and four referees for comments; and the Hewlett Foundation, the Joachim Herz Foundation, and the Deutsche Forschungsgemeinschaft through CRC TR 190 (project 280092119) for financial support. This paper was edited by Lance John Lochner.

Electronically published January 22, 2026

Journal of Political Economy, volume 134, number 3, March 2026.

© 2026 The University of Chicago. All rights reserved, including rights for text and data mining and training of artificial intelligence technologies or similar technologies. Published by The University of Chicago Press.
<https://doi.org/10.1086/739330>

in the United States has increased from 3.5% to over 23% between 1950 and 2020 (BEA 2021a), while the number of industrial robots per thousand workers has risen from 0.38 in 1993 to about 1.8 in 2017 (IFR 2018; BEA 2021b). There is growing evidence that these technologies have not only automated a range of tasks previously performed by workers and impacted the wage structure¹ but also led to polarization of employment and wages, meaning that the negative effects have concentrated on employment and wages in the middle of the wage distribution.² This pattern is intimately linked to the fact that many of the tasks that have been automated used to be performed by middle-skill workers.

There is no widespread agreement on why automation has been associated with polarization, however. One explanation, suggested by Autor (2014, 2015), is related to Polanyi's paradox, as captured by Polanyi's (1966, 4) statement that "we can know more than we can tell." Put simply, many of the manual and abstract tasks embed rich tacit knowledge, making them nonroutine. Because routine tasks are technologically easier to automate and are performed by middle-skill workers located in the middle of the wage distribution, new automation technologies have displaced labor from middle-skill occupations and have had their most negative effects on middle-pay worker groups.

In this paper, we provide an alternative complementary explanation: automation has focused on middle-skill tasks because these are the most profitable ones to automate. Specifically, low-skill tasks can be performed at lower labor expenses, reducing the cost advantage of machines relative to humans.

To develop this point, we build an assignment model, combining elements from the seminal contributions of Tinbergen (1956), Sattinger (1975), and Teulings (1995, 2005) together with the model of tasks and automation in Acemoglu and Restrepo (2022). Workers are distinguished by a single-dimensional skill index, which is distributed over an interval normalized to $[0, 1]$, and tasks are also distributed over the unit interval. For expositional ease, we refer to higher-index tasks as more "complex."³ Following

¹ On the spread of automation technologies over the past 80 years, see Autor (2015), Ford (2015), Graetz and Michaels (2018), Acemoglu and Restrepo (2020), and Acemoglu and Johnson (2023).

² The seminal contribution on the inequality and polarization effects of automation is Autor, Levy, and Murnane (2003). For a recent study of the effects of automation technologies on US wage inequality, see Acemoglu and Restrepo (2022). Employment polarization from automation is also documented in Goos, Manning, and Salomons (2009), Acemoglu and Autor (2011), and Autor and Dorn (2013). Acemoglu and Autor (2011) and Acemoglu and Restrepo (2022) provide evidence for wage polarization.

³ We show below more realistic configurations, where manual tasks that require skills that machines and algorithms do not currently fully possess can be incorporated into the model and can still be mapped into our single-dimensional tasks distribution.

the assignment literature, we assume that high-type workers have a comparative advantage in more complex tasks. Without automation, our model is identical to those in the previous literature and generates a monotone assignment pattern, with higher-skill workers performing more complex tasks. The distinguishing feature of our framework is that some tasks can be assigned to capital.

Under the assumption that capital does not have a comparative advantage for the most complex tasks and some additional restrictions on capital productivity, we prove that the equilibrium will take one of two forms: (1) *interior automation*, where capital performs a set of intermediate tasks; or (2) *low-skill automation*, where capital takes over all tasks below a certain threshold.⁴

Interior automation is the configuration that leads to polarization, and we provide conditions under which automation is indeed interior. These conditions depend on the comparative advantage of low-skill workers relative to capital, the cost of capital, and supplies of different types of labor, which together determine the equilibrium wage distribution with and without automation. Intuitively, when equilibrium wages (without automation) are sufficiently low for low-skill workers, tasks in the bottom of the complexity distribution are very cheap, and this reduces the profitability of performing them by capital. When this is the case, we also show that further automation leads to both wage and employment polarization. Therefore, in our model, polarization is closely linked to the fact that wages are already low at the bottom.

In addition to establishing the existence of a unique competitive equilibrium and characterizing the conditions under which capital takes over tasks from the middle of the skill distribution, we provide a series of comparative static results for marginal (local) and large (global) changes in automation.

Our first result clarifies the conditions under which automation is interior. In the baseline model, interior automation requires that wages at the bottom be sufficiently low relative to the productivity of low-skill workers and the effective cost of capital.⁵ To further highlight the role of wages at the bottom, we also consider a version of our model with a minimum wage. In this case, interior automation requires that the minimum wage is not too high; otherwise, low-skill labor is too expensive, and this induces low-skill automation. We complement this result by showing that a reduction

⁴ A third possibility is no automation, which is not interesting given our focus here and will be ruled out by assuming that capital productivity is sufficiently high to make some automation profitable in equilibrium.

⁵ This is the sense in which our explanation is complementary to Autor's (2014, 2015) account: interior automation is also more likely to emerge when low-skill workers are more productive at low-complexity tasks relative to machines.

in the supply of skills at the bottom also raises low-skill wages and makes a transition to low-skill automation more likely.

Our second result shows how a further expansion of interior automation—for example, driven by a decline in the price of capital goods—generates employment and wage polarization. Employment polarization here simply means that human workers are squeezed into smaller sets of tasks at the bottom and the top. Wage polarization takes a more specific form: relative wage changes increase as a function of the distance between the task that a skill type performs and the boundaries of the set of automated tasks. As a result, skill premia increase among worker types performing more complex tasks than those that are automated and decrease among worker types performing less complex tasks than the automated ones. Put differently, interior automation hurts (relatively) workers who are closer to the set of automated tasks. This is intuitive in view of the fact that workers closer to this set used to have a stronger comparative advantage for tasks that are now automated.

Third, we characterize the effects of automation on the level of real wages for different types of workers. As in Acemoglu and Restrepo (2022), whether the real wage of a given skill group declines depends on competing displacement and productivity effects, though in this paper we will provide more explicit conditions. One noteworthy result in this context is that the larger is the initial set of tasks that are automated, the more likely are the real wages of all skill types to increase. Moreover, we show that the productivity gains from automation are convex in the price of capital goods. This implies that as capital good prices (including costs of algorithmic automation) decline further, the productivity effect strengthens, ultimately eliminating negative wage impacts. These results together imply that the most negative consequences of displacement on workers will be at the early stages of the automation process.⁶

We present one more noteworthy result on wages, related to what we call Wiener's conjecture, after Wiener's (1950) pioneering study of automation. Wiener (1950, 189) claimed, "The automatic machine . . . is the precise economic equivalent of slave labor. Any labor which competes with slave labor must accept the economic consequences of slave labor." This conjecture does not seem to have been fully borne out by economic models or the data. On the theory side, Zeira (1998) and Acemoglu and

⁶ Because production in our model exhibits constant returns to scale and the cost of capital is constant, average wages always increase following automation, and this structure is also important for the result that when sufficiently many tasks are automated, the effects on all wages are positive. The consequences of automation on wages are more negative when the cost of capital is increasing in the stock of capital or when we depart from constant returns to scale or competitive markets. See Moll, Rachel, and Restrepo (2022) and Acemoglu, Jensen, and Restrepo (2025).

Restrepo (2018b) show that real wages will increase in the long run following automation. On the empirical side, although the real wages of low-education groups have declined over the past 40 years, automation was also rapid in the 1950s and 1960s; during these decades, wages for all demographic groups increased robustly. Our analysis suggests that Wiener's conjecture needs to be refined: different workers have different skills, and even if automated machines are like slave labor, they do not perfectly compete against all kinds of labor. Building on this intuition, we show that automation always reduces the real wages of worker types whose productivity schedule over tasks is sufficiently close to capital's productivity schedule (if such worker types exist, but they may not).

Fourth, we use the model to study global—as opposed to local—effects of automation, which result when there are large declines in costs of capital goods. We show that as long as these changes keep us in the region of interior automation, their effects are qualitatively the same as those of local changes. Ultimately, however, automation expands from the interior of the set of tasks to take over all low-skill tasks. When this happens, the pattern of polarization reverses. While an expansion in interior automation hurts workers in the middle of the skill distribution the most (and lowest-skill workers are to some degree sheltered), a switch from interior to low-skill automation has its most adverse effects on lowest-skill workers. Hence, our model predicts that as automation proceeds, its inequality implications may become worse, not just quantitatively but also qualitatively.⁷

Finally, we undertake a preliminary quantitative evaluation of the effects of automation in our framework. We calibrate our model parameters to match the 1980 US wage distribution and the effects of automation on different parts of the wage distribution between 1980 and 2016–17, as estimated in Acemoglu and Restrepo (2022). We show that our model matches these and a number of other moments in the US data quite well. We then consider (1) the impact of a further wave of automation (driven by a decline in the cost of capital of the same magnitude as in 1980–2016); (2) the implications of advances in artificial intelligence (AI), modeled as improvements in capital productivity in tasks that were previously not automated (as opposed to an across-the-board improvement in capital productivity or decline in the cost of capital); and (3) the effects of a minimum wage of \$16 an hour. Our analysis suggests that a further wave of automation similar to that of 1980–2016 would increase inequality by even more, because this change would have a bigger impact at the bottom of the wage distribution, given the patterns of comparative advantage implied by 1980–2016 data. AI is predicted to raise inequality by even

⁷ If in this process capital productivity in already automated tasks increases (which we refer to as deepening of automation), negative wage level effects become less likely.

more, because it expands the set of automated tasks significantly but does not induce as much deepening of automation—meaning productivity improvements in already automated tasks—as does a uniform fall in the cost of capital, which tends to increase productivity, benefiting labor of all types. Consequently, low-skill workers are harmed more by AI, while the highest-skill workers continue to benefit, because there is an increase in task services that are complementary to their skills. A sizable minimum wage like \$16 an hour for the United States, on the other hand, compresses the wage distribution considerably, raising wages at the bottom and reducing them at the top. The wage declines at the top are due to the fact that the minimum wage reduces overall employment and output, which then depresses demand for all tasks.

Our paper is related to several contributions in both the assignment and the automation literatures. In the assignment literature, early contributions include Timbergen (1956), Rosen (1974), Sattinger (1975, 1993), and Heckman and Sedlacek (1985). Our model builds more closely on Teulings (1995, 2005), Teulings and Gautier (2004), Costinot and Vogel (2010), and Stokey (2018). More recently, Lindenlaub (2017) extends assignment models to settings where workers and jobs have multidimensional characteristics and derives comparative statics with respect to the degree of complementarity between skills and manual and cognitive skill requirements of jobs. The major difference between all of these papers and our work is the presence of capital that can take over some tasks, which allows for an analysis of automation. From a technical point of view, these papers impose log supermodularity between all factors and job types, which turns the problem into one of monotone assignment. Our analysis of automation relaxes log supermodularity between capital, labor, and tasks (though for simplicity, we maintain log supermodularity between worker types and tasks).

In the automation literature, we build on earlier models where capital displaces workers in some of the tasks they used to perform. This literature and task-based models started with Zeira's (1998) seminal theoretical work and Autor, Levy, and Murnane's (2003) empirical study of the polarization and inequality effects of automation. Zeira's model includes only one type of labor and does not focus on inequality implications of automation. Many subsequent works—including Acemoglu and Zilibotti (2001), Acemoglu and Restrepo (2018a, 2018b), Berg, Buffie, and Zanna (2018), Jackson and Kanik (2020), Jaimovich et al. (2021) and Hémous and Olsen (2022)—allow only two types of workers, making it impossible to study wage polarization. Acemoglu and Autor (2011) study an economy with three types of workers and establish polarization when automation affects the middle group, but this structure does not allow a comprehensive analysis of the implications of different stages of automation on employment and wage patterns. Feng and Graetz (2020), Loebbing (2022), and Ales

et al. (2024) study task-based models with a continuum of labor types, though under more restrictive assumptions regarding comparative advantage between capital and labor. In particular, Feng and Graetz (2020) impose that automation is always interior, while Loebbing (2022) focuses on the case in which automation is always low skill. This contrasts with our focus, which is to understand when automation is interior and derive the conditions under which there is a transition to low-skill automation. Ales et al. (2024) study firm-level automation by imposing monotone comparative advantage of labor relative to capital in more complex tasks but also add machine indivisibility, which leads to a scale-dependent pattern of automation. At low scales, tasks of medium complexity are automated, whereas at high scales, indivisibilities become less relevant and automation is low skill. None of these papers contain our main characterization and comparative static results.

Acemoglu and Restrepo (2022) develop a general framework with multiple industries and multiple worker types to study the inequality effects of automation. In addition to providing empirical estimates of the effects of automation on US wage inequality, Acemoglu and Restrepo (2022) present a theoretical analysis of the implications of automation. Because their study lacks the specific structure imposed here (with one-dimensional heterogeneity on both the worker and the job complexity side and comparative advantage between workers and tasks), it does not contain results on which tasks capital will take over. Rather, they provide equations that specify how wages of different groups will be affected as a function of the total set of tasks that are automated and the ripple effects, which capture how different skill groups compete over marginal tasks. These ripple effects cannot be explicitly characterized given their assumptions and are studied empirically. Ocampo (2022) studies the assignment of a discrete set of workers and capital to a continuum of tasks. In his framework, matching is one to many, and occupations emerge endogenously as bundles of tasks assigned to different workers. Like Acemoglu and Restrepo (2022), Ocampo (2022) does not impose enough structure to characterize the pattern of automation and the equilibrium comparative statics as functions of primitives. In contrast to these works, our analysis enables a full characterization of equilibrium and its comparative statics.

The rest of the paper is organized as follows. Section II presents our baseline environment and defines a competitive equilibrium. Section III first establishes existence and uniqueness of equilibrium and some basic characterization results and then studies the conditions under which automation is interior. Section IV presents our main comparative static results for small changes in the cost of capital goods and derives the employment and wage polarization implications of automation. Section V considers global changes in automation technology and their equilibrium consequences. Section VI presents our quantitative analysis and studies

the consequences of various counterfactual technological and institutional changes. Section VII concludes. Appendix A includes several of the proofs omitted from the text, while appendix B (available online) contains a few additional proofs and information on data and computational methods.

II. Model

In this section, we introduce the basic economic environment, describe some of our assumptions and their motivations, and define a competitive equilibrium.

A. Environment

We consider a static economy with a unique final good, Y , produced from a continuum of workers with skills $s \in [0, 1]$ and a continuum of tasks $x \in [0, 1]$. The production of the final good is given as a constant elasticity of substitution aggregate of tasks:

$$Y = \left[\int_0^1 Y_x^{(\lambda-1)/\lambda} dx \right]^{\lambda/(\lambda-1)}, \quad (1)$$

where Y_x is the amount of task x and $\lambda > 0$ is the elasticity of substitution.⁸

All labor types are inelastically supplied, with a density function of $l: [0, 1] \rightarrow \mathbb{R}_{++}$ (which specifies the total endowment of each type of labor), and we assume that this density is continuous. We also assume that capital is produced out of final good with marginal cost $1/q$. We identify increases in q with greater capital productivity or equivalently lower prices of capital goods.

The task production functions are given by

$$Y_x = \int_0^1 \psi_{s,x} L_{s,x} ds + \psi_{k,x} K_x \quad (2)$$

for all $x \in [0, 1]$, where $\psi_{s,x} > 0$ and $\psi_{k,x} > 0$ denote the productivities of different factors in task x and $L_{s,x}$ and K_x are, respectively, the amounts of labor of type s and capital allocated to the production of task x . We assume that the factor productivities $\psi_{s,x}$ and $\psi_{k,x}$ are twice continuously differentiable. Labor market clearing requires

⁸ This production function imposes that all tasks have the same productivity/importance. This is without loss of generality, since we allow general task and factor-specific productivity functions $\psi_{s,x}$ and $\psi_{k,x}$ below, and any task-level differences can be subsumed into these.

$$\int_0^1 L_{s,x} dx = l, \text{ for all } s \in [0, 1], \quad (3)$$

and net output is

$$NY = Y - \frac{1}{q} \bar{K},$$

where $\bar{K} = \int_0^1 K_x dx$ is the aggregate capital stock and thus \bar{K}/q is total capital expenditure. Net output is also equal to consumption in this economy.

B. Competitive Equilibrium

An allocation in this economy is given by a collection of density functions, $L = \{L_s\}_{s=0}^1$, where $L_s : [0, 1] \rightarrow \mathbb{R}_+$ for each $s \in [0, 1]$ and a capital allocation $K : [0, 1] \rightarrow \mathbb{R}_+$. The density functions allocate labor supply of each type of labor to tasks, and the capital allocation function determines how much capital will be allocated to each task. This definition already incorporates nonnegativity constraints for all factors in all tasks. We describe an allocation with the shorthand $\{L, K\}$. For all $s \in [0, 1]$, we define the set $X_s : \{x \mid L_{s,x} > 0\}$ as the set of tasks performed by labor type s in this allocation and $X_k = \{x \mid K_x > 0\}$. We also use the terminology that tasks in the set X_k are *automated*.⁹

We additionally designate two price functions: first, a wage function $w : [0, 1] \rightarrow \mathbb{R}_+$, which gives the wage level, w_s , for each type of labor $s \in [0, 1]$; and second, a task price function $p : [0, 1] \rightarrow \mathbb{R}_+$, which determines the price p_x of each task $x \in [0, 1]$.

A (competitive) equilibrium is defined as an allocation $\{L, K\}$ and price functions $\{w, p\}$ such that final good producers maximize profits, taking task prices as given; task producers maximize profits, taking task and factor prices as given; and all markets clear. Task producers' profit maximization implies that wages equal marginal products of the relevant labor types, that is,

$$\begin{aligned} w_s &= p_x \psi_{s,x} \text{ for all } x \in X_s, \\ w_s &\geq p_x \psi_{s,x} \text{ for all } x \in [0, 1], \end{aligned} \quad (4)$$

while the cost of capital must be equal to the marginal product of capital, that is,

⁹ A more stringent definition of automated tasks might additionally require that these tasks are not simultaneously performed by labor, i.e., $L_{s,x} = 0$ for all s . Under our assumptions 2 and 3, the set of tasks that are performed by both labor and capital in equilibrium is of measure zero, and hence this distinction is not relevant.

$$\begin{aligned} \frac{1}{q} &= p_x \psi_{k,x} \quad \forall x \in X_k, \\ \frac{1}{q} &\geq p_x \psi_{k,x} \quad \forall x \in [0, 1]. \end{aligned} \tag{5}$$

Final good producers' profit maximization in turn implies that task prices equal the marginal products of tasks in final good production, that is,

$$p_x = \left(\frac{Y}{Y_x} \right)^{1/\lambda} \text{ for all } x \in [0, 1]. \tag{6}$$

We note that the first welfare theorem holds in our model and the equilibrium allocation maximizes net output (consumption) subject to labor market clearing (3).

C. Assumptions and Motivation

We now describe some of the assumptions we will use in our main analysis.

Since the economy exhibits constant returns to scale and can produce the final good linearly from capital and capital from the final good, in principle its output may be unbounded. Our first assumption ensures that the cost of capital is not so low as to generate infinite output.

ASSUMPTION 1 (Bounded output). The cost of capital satisfies

$$\frac{1}{q} > \frac{1}{q_\infty} = \left(\int_0^1 \psi_{k,x}^{\lambda-1} dx \right)^{1/(\lambda-1)}$$

(where the lower bound $1/q_\infty$ is derived as the marginal product of capital if all tasks are performed by capital, using eqq. [1], [5], and [6]).

Our second assumption follows the assignment literature (e.g., Teulings 1995, 2005; Costinot and Vogel 2010) and imposes *comparative advantage* (among workers). This means, in particular, that the productivity advantage of higher-skilled workers increases more than proportionately with the task index among workers. We impose this assumption both to simplify the analysis and also to maximize the similarity of our benchmark environment to the previous literature, which will clarify that all of the new results here are driven by the automation margin.

ASSUMPTION 2 (Comparative advantage among workers). For all $s > s'$, we have $\psi_{s,x}/\psi_{s,x'} > \psi_{s',x}/\psi_{s',x'}$ for all $x > x'$.

Comparative advantage ensures that without capital, the equilibrium will assign higher-skilled workers to higher-indexed tasks.¹⁰

We next present a motivating example that provides a simple illustration of comparative advantage. This example has the additional benefit of showing how multidimensional skills can be mapped into our setup with a one-dimensional skill index.

EXAMPLE 1. Suppose that each task $x \in [0, 1]$ involves a combination of abstract and manual activities. Specifically, the productivity of a worker with skill level $s \in [0, 1]$ in task x will be a function of this worker's abstract and manual skills, denoted by the vector (a_s, m_s) :

$$\psi_{s,x} = a_s^x m_s^{1-x}.$$

In the context of this example, a sufficient condition for comparative advantage is for workers' skill endowments (a_s, m_s) to satisfy

$$\frac{a_s}{m_s} > \frac{a_{s'}}{m_{s'}} \text{ for all } s > s'.$$

To see this, let

$$\Lambda = \log \frac{\psi(x, s)}{\psi(x, s')} = \log \frac{a_s^x m_s^{1-x}}{a_{s'}^x m_{s'}^{1-x}},$$

and by assumption, we have $\partial\Lambda/\partial x = [\log a_s - \log m_s] - [\log a_{s'} - \log m_{s'}] > 0$. A sufficient condition for absolute advantage is for a_s to be strictly increasing and m_s to be nondecreasing in s .

Motivated by this pattern of comparative advantage, we also refer to higher-index tasks as *more complex tasks*, as in Teulings (1995, 2005).

The other key dimension of our model concerns the productivity of capital relative to different labor types. Crucially, here we do not assume supermodularity. However, it is convenient to put sufficient structure on the comparative advantage of capital to have a simple characterization of the set of tasks, X_k , that are assigned to capital. The next assumption achieves this.

ASSUMPTION 3 (Comparative advantage of capital). For all $s \in [0, 1]$, $\psi_{k,x}/\psi_{s,x}$ is quasi-concave in x .

¹⁰ Since we refer to higher levels of the skill index s as more skilled, it is natural to presume that wages should increase in s . A simple way to ensure this would be to impose absolute advantage, i.e., $\psi_{s,x} > \psi_{s',x}$ for all $s > s'$ and all tasks x , but we do not formally impose this restriction because it is not needed for the rest of our analysis.

This assumption rules out situations in which the direction of comparative advantage for capital changes more than once for any given level of skill. Put differently, assumption 3 allows some skill types to have comparative advantage in lower-index tasks relative to capital and then again in higher-index tasks after a certain threshold. But it rules out more than one such switch. We prove in proposition 2 that this is necessary and sufficient for the set X_k of tasks assigned to capital to be convex. We adopt assumption 3 in the text for expositional simplicity. In appendix B, we extend our main characterization result (proposition 3) to the case where assumption 3 is relaxed.

We next illustrate this assumption with the environment considered in example 1.

EXAMPLE 1 (Continued). Assumption 3 is ensured in this example when $\psi_{k,x}$ is log concave in x . Since the productivity of each labor type is log linear in x , log concavity of capital productivity implies quasi-concavity of all relative productivity schedules $\psi_{k,x}/\psi_{s,x}$.

III. Characterization of Equilibrium and Interior Automation

In this section, we establish existence and uniqueness of a competitive equilibrium and study the conditions under which tasks from the middle of the skill distribution are automated.

A. Existence and Uniqueness

PROPOSITION 1 (Existence and uniqueness). Suppose that assumption 1 holds. Then, a competitive equilibrium always exists and is essentially unique in the sense that wage and price functions are uniquely determined. If in addition assumptions 2 and 3 hold, the competitive equilibrium is unique.

Existence follows from the fact that the competitive equilibrium maximizes net output, which is a continuous function of the allocation. Essential uniqueness, on the other hand, is a consequence of the fact that net output is a concave function of the allocation. The reason why the competitive equilibrium is essentially unique—but not fully unique without assumptions 2 and 3—is straightforward to see: some tasks may be produced at the same cost using different factors, creating indeterminacy of equilibrium allocations. Assumptions 2 and 3 rule out such indeterminacy: assumption 2 imposes strict comparative advantage between any two types of labor and, together with assumption 3, implies that on any subset of tasks of positive measure, there can be at most one

type of labor with a productivity schedule parallel to capital's productivity schedule.¹¹

B. Interior Automation

The next proposition confirms that, as mentioned above, assumption 3 is sufficient to ensure that the set of tasks allocated to capital, X_k , is convex.

PROPOSITION 2 (Convexity of assignment). Suppose that assumptions 1 and 2 hold.

1. The allocation of labor across tasks is monotone, meaning that for any $s > s'$, if $L_{s,x} > 0$, then $L_{s',x'} = 0$ for all $x' \geq x$.
2. The set of automated tasks in equilibrium, X_k , is convex for all labor endowment functions and capital productivity levels if and only if assumption 3 holds.

The first part of this proposition confirms that the monotonicity obtained in assignment models with log supermodularity continues to hold in our model. The second part implies that we can focus on a convex set of automated tasks. A convex set of automated tasks leaves four feasible configurations:

1. No automation: where $X_k = \emptyset$ (because the cost of capital is too high).
2. Interior automation: where $X_k = [\underline{x}, \bar{x}]$, with $0 < \underline{x} < \bar{x} < 1$, and thus both the most complex and the least complex tasks are assigned to some labor types.¹²
3. Low-skill automation: where $X_k = [0, \bar{x}]$, with $0 < \bar{x} < 1$, and thus all tasks below a certain threshold of complexity are automated.
4. High-skill automation: where $X_k = [\underline{x}, 1]$, with $0 < \underline{x} < 1$, such that all tasks above a certain complexity threshold are automated.

We refer to the third configuration as low-skill automation, since capital takes over tasks that used to be performed by lower-skilled workers (as in the first part of proposition 2), and analogously we refer to the fourth case as high-skill automation. The first case—no automation—is not of great

¹¹ If a single type of labor has a productivity schedule that is parallel to capital's productivity schedule on the set of automated tasks, this does not create any indeterminacy in allocations because a single type of labor has no mass. Formally, an allocation is a collection of densities from an L^p space where any two densities that are equal almost everywhere are identified and represent the same allocation.

¹² In this statement, we impose that the boundary tasks \underline{x} and \bar{x} are performed by capital. This is for notational simplicity and is without loss of any generality.

interest given the focus of this paper, and we assume below that the cost of capital is sufficiently low so as to ensure automation.

We next study the conditions under which automation will be interior. The least complex task, $x = 0$, is cheaper to produce by the least skilled worker type, $s = 0$, than by capital if

$$\frac{w_0}{\psi_{0,0}} < \frac{1/q}{\psi_{k,0}}. \quad (7)$$

When inequality (7) is satisfied, we cannot have low-skill automation, and an analogous condition rules out high-skill automation. Hence, from proposition 2, automation must be interior. Condition (7) is intuitive. It requires that the effective wage of the least skilled workers in the least complex task (wage divided by productivity) is less than the effective cost of capital in that task (the cost of capital, $1/q$, divided by the productivity of capital in that task). Whether this condition is satisfied depends on the shape of comparative advantage schedules ψ , capital productivity q , and the labor supply profile l (which jointly determine the equilibrium wage for the least skilled type, w_0). The following two conditions allow us to characterize the equilibrium both in cases where inequality (7) is satisfied and in cases where it is not.

CONDITION 1 (Local comparative advantage of capital).

1. The least skilled workers have local comparative advantage relative to capital in the least complex tasks, that is,

$$\frac{\partial \log \psi_{0,0}}{\partial x} < \frac{\partial \log \psi_{k,0}}{\partial x}.$$

2. The most skilled workers have local comparative advantage relative to capital in the most complex tasks, that is,

$$\frac{\partial \log \psi_{1,1}}{\partial x} > \frac{\partial \log \psi_{k,1}}{\partial x}.$$

CONDITION 2 (Global comparative advantage of capital).

1. The least skilled workers have global comparative advantage relative to capital in the most complex tasks, that is,

$$\frac{\psi_{0,0}}{\psi_{k,0}} < \frac{\psi_{0,1}}{\psi_{k,1}}.$$

2. The most skilled workers have global comparative advantage relative to capital in the most complex tasks, that is,

$$\frac{\psi_{1,0}}{\psi_{k,0}} < \frac{\psi_{1,1}}{\psi_{k,1}}.$$

Condition 1 imposes only local comparative advantage: comparing the least (most) skilled workers with capital around the least (most) complex tasks. This comparative advantage pattern does not have to hold globally. Condition 2 in contrast imposes global comparative advantage: comparing the least and most skilled workers to capital in the least versus most complex tasks.

PROPOSITION 3 (Interior automation). Suppose that assumptions 1–3 and condition 1 hold. Then, there exist thresholds $q_0 < q_\infty$ and $q_m \in (q_0, q_\infty]$ such that:

1. For $q \leq q_0$, there is no automation, that is, $X_k = \emptyset$.
2. For $q \in (q_0, q_m)$, automation is interior, that is, $X_k = [\underline{x}, \bar{x}]$, with $0 < \underline{x} < \bar{x} < 1$. Additionally:
3. If condition 2 holds, then $q_m < q_\infty$, and for $q \geq q_m$, automation is low skill, that is, $X_k = [0, \bar{x}]$, with $0 < \bar{x} < 1$.
4. If condition 2.1 is strictly violated but condition 2.2 holds, then $q_m = q_\infty$ such that automation is interior for all $q \in (q_0, q_\infty)$.
5. If conditions 2.1 and 2.2 are both strictly violated, then $q_m < q_\infty$, and for $q \geq q_m$, automation is high skill, that is, $X_k = [\underline{x}, 1]$, with $0 < \underline{x} < 1$.¹³

If assumptions 1–3 hold but condition 1 does not, then automation is not interior.

Proposition 3 is our first main result and provides a complete characterization of the different patterns of automation that can arise in our model. The first part of the proposition establishes that capital productivity has to cross some threshold level q_0 to induce any automation. The second part is the most important result of the proposition. It establishes that when capital productivity crosses this threshold q_0 and assumptions 1–3 and condition 1 hold, automation always starts in the interior.

The next three parts show that when capital productivity increases further, three different scenarios are possible. First, if condition 2 holds—which means that both the least and the most skilled workers have global

¹³ Note that a configuration where condition 2.1 holds while condition 2.2 is violated is not compatible with the comparative advantage among workers imposed in assumption 2. In knife-edge cases where one of the two conditions is violated weakly, meaning that the respective inequality is replaced by equality, additional information about comparative advantage schedules is needed to determine the pattern of automation, and we omit those cases to save space.

comparative advantage relative to capital in the most versus the least complex tasks—then automation transitions from interior to low skill at some threshold level of capital productivity, denoted by q_m (and characterized in app. sec. A2.3). Second, if only the most skilled workers have global comparative advantage relative to capital in the most complex tasks, then automation remains interior indefinitely. Third, if both the least skilled and the most skilled workers have global comparative advantage relative to capital in the least complex tasks, then automation transitions from interior to high skill after capital productivity exceeds the threshold.

Finally, the last result in the proposition shows that condition 1 is necessary for interior automation. This part of the proposition is developed further in proposition A1 in appendix section A1.

Figure 1 shows the assignment of tasks to labor and capital in the case of interior automation. Tasks in the set $X_k = [\underline{x}, \bar{x}]$ are assigned to capital, while the remaining tasks are performed by the labor types indicated on the vertical axis.

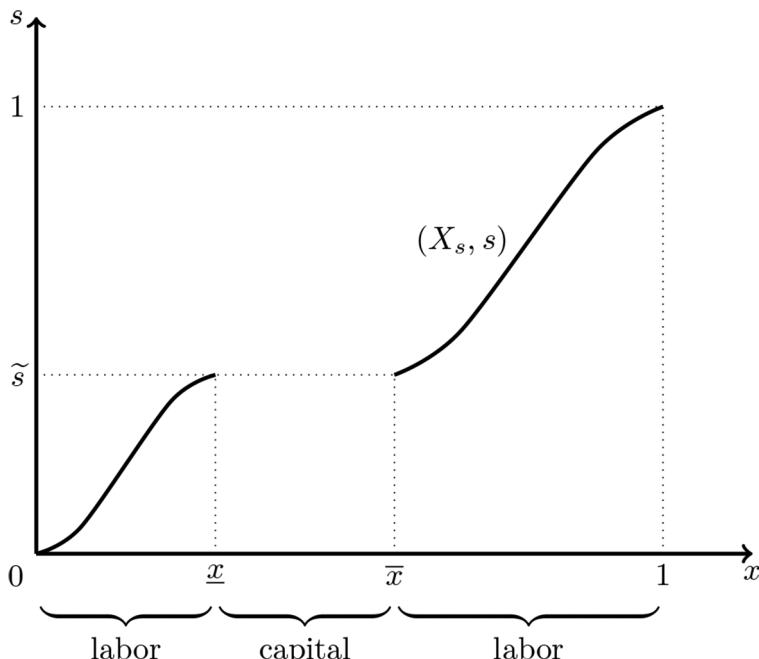


FIG. 1.—Assignment of tasks to capital and labor. Tasks in the set $X_k = [\underline{x}, \bar{x}]$ are assigned to capital, while the remaining tasks are assigned to the labor types indicated by the graph (X_s, s) . In particular, tasks $x < \underline{x}$ are assigned to worker types $s < \bar{s}$, while tasks $x > \bar{x}$ are assigned to labor types $s > \tilde{s}$.

Inequality (7) provides further intuition for why automation may affect middle-skill occupations most. Fixing $\psi_{0,0}$ and treating the wage of the least skilled worker, w_0 , parametrically, there are two ways in which this inequality is satisfied: either $1/(q\psi_{k,0})$ is high or w_0 is low. The first captures the economic forces proposed by Autor (2014, 2015): many of the tasks performed by lower-skill workers may be hard to automate because they require a combination of tacit knowledge and manual dexterity. The second is what we have emphasized in the introduction: wages at the bottom are too low to make automation economically profitable.

Proposition 3 clarifies that these two explanations are linked because the wage is endogenous. However, they are also distinct, and one way of illustrating this is to consider variations in the wages at the bottom of the distribution, holding the other parameters of the model constant. The simplest way of doing this is by imposing a minimum wage in the model, which we discuss briefly in proposition 4. In the presence of a binding minimum wage w , the equilibrium involves rationing: some worker types may not be hired. This requires an obvious change in the definition of equilibrium, which we omit to save space. It is also straightforward to see that the set of rationed workers will always be of the form $[0, \underline{s}]$ (see Teulings 2000). Except for rationing, the same equilibrium conditions as in our analysis so far apply. Then we have:

PROPOSITION 4 (Minimum wages and automation). Suppose that assumptions 1–3 and condition 1 hold and let $q \in (q_0, q_m)$ so that in the competitive equilibrium without the minimum wage, we have interior automation. Now consider a minimum wage of $w > 0$, which leads to the rationing of workers with skills in $[0, \underline{s}]$. If in addition we have

$$\frac{\partial \log \psi_{2,0}}{\partial x} \geq \frac{\partial \log \psi_{k,0}}{\partial x}, \quad (8)$$

then inequality (7) is violated, and we transition to low-skill automation.

Intuitively, without the minimum wage, labor performing low-skill tasks tends to be cheap, and this makes automating these tasks unprofitable, ensuring interior automation. When there is a binding minimum wage, skills at the bottom of the distribution become more expensive, and this increases the profitability of automating some of the tasks previously performed by low-skill workers (and also causes unemployment). It is also useful to observe the role of condition (8): without this condition, some of the workers with skill above \underline{s} may find it profitable to take the lowest-complexity tasks.¹⁴ The implication that higher wages for low-skill labor

¹⁴ This condition is compatible with condition 1, since the comparison is for different skill levels, and the juxtaposition of these two conditions highlights that in our model, conditions

lead to more automation of the tasks performed by impacted workers is consistent with the empirical findings of Hémous et al. (2025), who exploit cross-country and cross-firm variation in low-skill wages and trace their effects on automation innovations using patent keywords.

An alternative way to increase wages at the bottom of the distribution is to reduce the labor supply of low-skill workers relative to high-skill workers. We consider such a change in labor supply in the following proposition.

PROPOSITION 5 (Labor supply and automation). Suppose that assumptions 1–3 as well as conditions 1 and 2 hold so that automation transitions from interior to low skill at the threshold $q_m \in (q_0, q_\infty)$. Consider a change in labor supply such that $\Delta \log l_s < \Delta \log l_{s'}$ for all $s < s'$ (an increase in the relative supply of more skilled workers). Then, the threshold q_m declines, that is, $\Delta q_m < 0$. Thus, if $q \in (q_m + \Delta q_m, q_m)$, automation transitions from interior to low skill in response to the labor supply change.

The proposition shows that an increase in the relative supply of high-skill workers can induce a transition from interior to low-skill automation. This confirms our intuition about the critical role of low-skill workers' wages: an increase in relative skill supply renders low-skill workers scarce and raises their wages, which then makes it more profitable to automate their jobs. The implications of proposition 5 are consistent with the empirical findings of Clemens, Lewis, and Postel (2018), who exploit the end of the Bracero Program that led to the exclusion of about half a million low-skill Mexican farmworkers from the US agricultural sector and show that this led to the substitution of capital for the labor of these workers (which is equivalent to low-skill automation in our framework). We further discuss the implications of labor supply changes for wage inequality in proposition 12.

C. Characterization of Equilibrium

We next provide a characterization of equilibrium when assumptions 1–3 hold and $q \geq q_0$ (so that there is automation in equilibrium). Under these assumptions, the set of automated tasks takes the form $X_k = [\underline{x}, \bar{x}]$, with $\underline{x} \leq \bar{x} \leq 1$. This leaves the sets $[0, \underline{x})$ and $(\bar{x}, 1]$ for labor, with workers with skills above a threshold $\tilde{s} \in [0, 1]$ employed in $(\bar{x}, 1]$ and those below \tilde{s} employed in $[0, \underline{x})$.

By standard arguments from the assignment literature, the allocation of skills to tasks in either of the two sets $[0, \underline{x})$ and $(\bar{x}, 1]$ can be described by an *assignment function* $X : [0, \tilde{s}) \cup (\tilde{s}, 1] \rightarrow [0, \underline{x}) \cup (\bar{x}, 1]$ mapping skills to

for the comparative advantage of capital relative to labor are endogenous: which skill type's productivity is compared with capital's productivity is determined in equilibrium.

tasks.¹⁵ This function is differentiable (except at the threshold \tilde{s}), strictly increasing and onto (see Costinot and Vogel 2010). These properties are illustrated in figure 1, except that we find it more intuitive to plot its inverse, X^{-1} , and thus have skills on the vertical axis and tasks on the horizontal axis. Consequently, when the assignment function X shifts up (down), its inverse X^{-1} (plotted in fig. 1) will shift down (up).

Condition (4) implies that every worker type is assigned to the task in which its marginal value product is maximized. Thus,

$$\log w_s = \log p_{X_s} + \log \psi_{s,X_s} = \max_x \{\log p_x + \log \psi_{s,x}\}.$$

An envelope argument then yields the differential equation

$$(\log w_s)' = \frac{\partial \log \psi_{s,X_s}}{\partial s} \quad \forall s \neq \tilde{s}, \quad (9)$$

which we can think of as determining wages given assignment and a boundary condition (where $(\log w_s)'$ denotes the derivative of the function $\log w_s$ with respect to s). When condition 1 holds and $q \in (q_0, q_m)$ such that automation is interior, the boundary condition is provided by the requirement that in both tasks \underline{x} and \bar{x} , production must be equally costly with capital and skill \tilde{s} :

$$\frac{w_{\tilde{s}}}{\psi_{\tilde{s},\underline{x}}} = \frac{1/q}{\psi_{k,\underline{x}}} \text{ and } \frac{w_{\tilde{s}}}{\psi_{\tilde{s},\bar{x}}} = \frac{1/q}{\psi_{k,\bar{x}}}.$$

When automation is low skill, we have $\underline{x} = 0$ and $\tilde{s} = 0$, and the first equality becomes an inequality— $w_{\tilde{s}}/\psi_{\tilde{s},\underline{x}} \geq 1/(q\psi_{k,\underline{x}})$ —so that it is weakly more costly to produce task $\underline{x} = 0$ with skill $\tilde{s} = 0$ than with capital. The second equality still provides a relevant boundary condition. So, taking logs, we have

$$\log w_{\tilde{s}} = \log \psi_{\tilde{s},\bar{x}} - \log \psi_{k,\bar{x}} - \log q \geq \log \psi_{\tilde{s},\underline{x}} - \log \psi_{k,\underline{x}} - \log q, \quad (10)$$

with equality when automation is interior and with $\tilde{s} = \underline{x} = 0$ when automation is low skill. An analogous condition can be derived for the case of high-skill automation. Note further that when automation is interior, the wage function characterized by (9) and (10) has a kink point at \tilde{s} , where the assignment function jumps upward.¹⁶

¹⁵ Recall that X_s was defined as the set of tasks performed by skill s , and thus in general X should be a correspondence. However, under assumptions 1–3, X_s is a singleton, and henceforth we treat X as a function.

¹⁶ The wage function is continuous at \tilde{s} because the productivity schedule $\psi_{s,x}$ is continuous in s . If there were an upward jump in wages at \tilde{s} , workers immediately below \tilde{s} would relocate to the tasks assigned to workers immediately above the threshold and increase their wage.

Intuitively, equation (9) ensures that all workers find it optimal to sort into the tasks assigned to them. This requires that the marginal return to skill at any level s is given by the marginal productivity gain in the task assigned to s , X_s .

Next, we can combine the equilibrium conditions for wages in (4), task prices (6), and task production (2) to obtain an expression for inverse labor demand:

$$w_s = Y^{1/\lambda} \psi_{s, X_s}^{(\lambda-1)/\lambda} L_{X_s}^{-1/\lambda},$$

where L_X is the marginal density of labor over tasks. A change of variable allows us to express this density in terms of the density of labor over skills, $X'_s L_{X_s} = l_s$, for $s \neq \tilde{s}$. Using this relationship and rearranging, we obtain the labor demand curve as

$$\frac{l_s}{X'_s} = \frac{Y \psi_{s, X_s}^{\lambda-1}}{w_s^\lambda} \quad \forall s \neq \tilde{s}. \quad (11)$$

The labor demand curve here takes the form of a differential equation for the assignment function, given wages. If automation is interior, the assignment function has two branches, one on $[0, \tilde{s})$ and one on $(\tilde{s}, 1]$. If automation is low skill instead, only the upper branch exists. For the lower branch, the boundary condition is

$$\lim_{s \nearrow \tilde{s}} X_s = \underline{x}, \quad (12)$$

whereas for the upper branch, the boundary condition is given by

$$\lim_{s \searrow \tilde{s}} X_s = \bar{x}. \quad (13)$$

Intuitively, if labor demand in task X_s is high (e.g., because aggregate output is high or the wage of skill s is low), equation (11) requires that the density of labor supplied to task X_s is high as well. This is achieved by a shallower slope for the assignment function X'_s , which means that more workers are squeezed into fewer tasks in the neighborhood of X_s .

Overall, we have a two-dimensional system of differential equations (and boundary conditions) for wages and assignment. This system fully characterizes equilibrium together with the production function (1), the capital allocation rule

$$K = \operatorname{argmax} \left\{ Y - \frac{\bar{K}}{q} \right\},$$

and the requirement that any task is assigned to some production factor. If automation is interior, this requires $X_0 = 0$ and $X_1 = 1$. If automation is low skill, only $X_1 = 1$ is required.

Our characterization displays the two channels via which automation (e.g., driven by a decline in the cost of capital/an increase in q) affects wages and assignment. The first is a *displacement effect*, as in Acemoglu and Restrepo (2022): automation reduces the boundary condition for wages (10) and hence, given assignment, the wage of worker type \tilde{s} . This implies, in particular, that workers directly competing with capital must either relocate to other tasks or accept a wage decline in proportion to the reduction of the cost of capital. The second is a *productivity effect*, driven by the fact that a lower cost of capital raises aggregate output Y . From equation (11), the productivity effect raises labor demand in all tasks proportionately and, for a given assignment, wages for all skill levels rise proportionately as well.

Finally, we can derive a simple expression for the share of capital in national income, which will be useful when discussing the productivity effects of automation. Combining the task production function (2), equation (5), and task prices (6) for the marginal product of capital, we obtain

$$\frac{1}{q} = Y^{1/\lambda} \psi_{k,x}^{(\lambda-1)/\lambda} K_x^{-1/\lambda}.$$

Then, solving for capital, integrating over $[x, \bar{x}]$, and dividing by qY yields the share of capital in gross output as

$$\alpha_k = \frac{\bar{K}/q}{Y} = \Gamma_k q^{\lambda-1}, \quad (14)$$

where

$$\Gamma_k = \int_x^{\bar{x}} \psi_{k,x}^{\lambda-1} dx$$

is the task share of capital, which is a productivity-weighted measure of the set of automated tasks. Equation (14) shows that a decline in the cost of capital has two distinct effects on the capital share: a capital deepening effect (as captured by $q^{\lambda-1}$), the sign of which depends on whether tasks are complements or substitutes; and the impact following from the expansion of the task share of capital, which is the counterpart of the displacement effect on wages discussed above.

IV. Local Effects of Automation

In this section, we suppose that assumptions 1–3 as well as condition 1 hold and that $q \in (q_0, q_m)$ so that automation is interior. We then study the implications of a small decline in the cost of capital goods (an increase in q), which expands the set of automated tasks. Our main results characterize the polarization and inequality consequences of automation.

A. Employment Polarization

PROPOSITION 6 (Automation and employment polarization). Suppose that assumptions 1–3 and condition 1 hold, let $q \in (q_0, q_m)$, and consider a small increase in the productivity of capital $d \log q > 0$. Then,

$$dx < 0 \text{ and } d\bar{x} > 0$$

(automation expands in both directions) and

$$dx_s < 0 \text{ for all } s \in (0, \tilde{s}) \text{ and } dx_s > 0 \text{ for all } s \in (\tilde{s}, 1)$$

(the assignment function shifts down below the set of automated tasks and shifts up above the set).

Moreover, if $\lambda \geq 1$, the labor share always decreases. If $\lambda < 1$, there exists a threshold for capital productivity $\hat{q} > q_0$ such that the labor share decreases if $q \in (q_0, \hat{q})$.

The first part of this proposition establishes that a (small) decline in the cost of capital goods (or an increase in the productivity of capital) always expands the set of automated tasks on both sides, and relatedly, it shifts the assignment of workers further toward the two extremes of the task distribution, as shown in figure 2. This result thus implies that the employment polarization pattern documented in Autor, Levy, and Murnane (2003), Goos, Manning, and Salomons (2009), and Acemoglu and Autor (2011) always applies so long as we consider a small increase in the set of automated tasks, starting from interior automation.¹⁷

One implication of this result is that some of the workers may end up performing more complex tasks after further automation (if $d\tilde{s} < 0$, the case shown in fig. 2). This is consistent with the empirical results presented in Dauth et al. (2021) showing how German firms retrain and promote blue-collar workers to more demanding technical tasks after the introduction of industrial robots.

The second part of the proposition provides conditions under which the labor share declines. There are two channels via which automation affects the labor share (see also our discussion of eq. [14] for the capital share). First, the expansion of the set of automated tasks, established in the first part of the proposition, always decreases the labor share. Second, productivity gains in tasks that are already automated (deepening of automation) decrease the labor share when tasks are substitutes ($\lambda \geq 1$) but raise it when tasks are complements ($\lambda < 1$). Yet even when tasks are

¹⁷ It is also straightforward to show that if there were technological constraints on what tasks could be automated (as in Acemoglu and Restrepo 2018b) and these dictated that only tasks in the set $[x, \bar{x}]$ could be automated, and we considered an expansion of the set with $dx < 0$ and $d\bar{x} > 0$, then the same employment polarization result would hold.

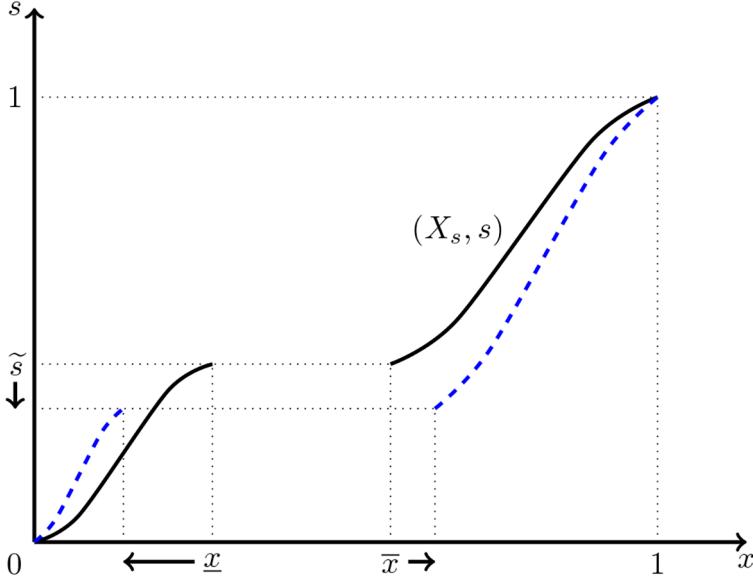


FIG. 2.—Employment polarization. In response to a small increase in capital productivity, the set of automated tasks expands in both directions, and workers move toward the extremes of the task distribution, here illustrated for the case with $d\tilde{s} < 0$ (dashed line). Recall that we are plotting the inverse of the assignment function, so when the assignment function shifts up, our plotted function shifts down.

complements, the proposition establishes that the expansion of the set of automated tasks dominates and the labor share declines in the initial stages of automation (when the productivity of capital is small).

B. Wage Polarization

The next proposition gives one of our most important results:

PROPOSITION 7 (Automation and wage polarization). Suppose that assumptions 1–3 and condition 1 hold, let $q \in (q_0, q_m)$, and consider a small increase in the productivity of capital $d \log q > 0$. Then, there is wage polarization in the sense that skill premia increase above the threshold task \tilde{s} and decrease below this threshold. Or equivalently,

$$\begin{aligned} d \log w_s &> d \log w_{s'} \text{ for all } s < s' \in (0, \tilde{s}], \\ d \log w_s &< d \log w_{s'} \text{ for all } s' > s \in [\tilde{s}, 1). \end{aligned}$$

The wage polarization result contained in proposition 7 is similar to the finding in Acemoglu and Autor (2011), but as discussed in the introduction, their result was a direct consequence of the fact that there were three types

of workers, and automation was assumed to affect the middle type. Here, we see that wage polarization reflects more general forces and applies throughout the distribution, regardless of exactly where automation is taking place (provided that we start from interior automation). We are not aware of other results of this sort in the literature.

The economics of this result is again related to the competing displacement and productivity effects. The displacement effect directly harms the earnings of workers who used to perform the previously automated tasks, while the productivity effect benefits all workers symmetrically. Notably, the displacement effect does not impact just directly affected workers (whose previous tasks are taken over by capital) but all workers, because of the general pattern of substitutability between worker types. These ripple effects are also present in Acemoglu and Restrepo (2022), but in our setting, they depend on only the distance of a skill group to the threshold type \tilde{s} . This, combined with the symmetric productivity effects, yields the result in proposition 7.

Proposition 7 establishes how skill premia change, generating a pattern of wage polarization. Other important questions are whether the real wage level of some worker types will decline following the expansion in automation and whether the top or the bottom of the wage distribution will be more heavily impacted. The next proposition answers these questions.

PROPOSITION 8 (Automation and wage levels). Suppose that assumptions 1–3 and condition 1 hold, let $q \in (q_0, q_m)$, and consider a small increase in the productivity of capital $d \log q > 0$.

1. The average wage in the economy always increases.
2. There exists a threshold for capital productivity $\hat{q} > q_0$ such that if $q \in (q_0, \hat{q})$, then for some $\delta_1, \delta_2 > 0$, we have $d \log w_s < 0$ for all $s \in (\tilde{s} - \delta_1, \tilde{s} + \delta_2)$.
3. Suppose that there exists s' such that $\psi_{s,x}/\psi_{k,x}$ is constant in x . Then for some $\delta_1, \delta_2 > 0$, we have $d \log w_s < 0$ for all $s \in (s' - \delta_1, s' + \delta_2)$.
4. Suppose that $\psi_{0,0}/\psi_{k,0} < \psi_{0,1}/\psi_{k,1}$. Then there exists a threshold for capital productivity $\tilde{q} < q_m$ such that if $q \in (\tilde{q}, q_m)$, the inequality between the top and the bottom of the skill space increases. That is,

$$d \log w_0 < d \log w_1.$$

The first part follows immediately from Euler's theorem given constant returns to scale, since in our economy net output equals the aggregate wage bill:

$$NY = \int_0^1 w_s l_s ds.$$

Because labor supply is unchanged, log differentiating this equation yields

$$\int_0^1 \frac{\alpha_s}{1 - \alpha_k} d \log w_s ds = d \log NY = \frac{\alpha_k}{1 - \alpha_k} d \log q > 0,$$

where α_s and α_k are the income shares of skill s and capital, respectively, and so the leftmost term is the change in the average wage in the economy, while the second equality is a direct implication of Hulten's theorem. Hence, the average wage always increases following an expansion in automation.¹⁸

The second part is also intuitive. When the initial level of capital productivity is low, the set of automated tasks is small. This implies that a marginal increase in q generates only a small productivity effect, and the most affected worker type, \tilde{s} , necessarily experiences a real wage decline (because of the displacement effect). In fact, the decline in the real wage extends to a set of workers around \tilde{s} because the wage effects are continuous in skills. This result highlights the importance of the magnitude of the productivity effect, which we characterize in section IV.C.

The third part provides a refinement of Wiener's conjecture, discussed in the introduction. Namely, if a worker type has a productivity profile very similar to that of capital, then Wiener's intuition that production using capital will cause the impoverishment of this worker type is correct.¹⁹ However, even though all labor types are competing against capital, wages will not fall for all workers but only for worker types whose overall comparative advantage is very similar to that of capital. Indeed, we know from the first part that average wages and hence the wages of some skill types have to increase (and, in fact, it is possible for all wages to increase).

Finally, the fourth part shows that automation widens the inequality between high- and low-skill workers, at least if the productivity of capital is high enough. The intuition for this result is as follows. If $\psi_{0,0}/\psi_{k,0} < \psi_{0,1}/\psi_{k,1}$, then automation will proceed in an unbalanced way, approaching the bottom of the task space as q grows (see sec. V). As automation tilts toward the bottom, so do its displacement effects on wages, reducing wages at the bottom relative to the top of the skill space.

¹⁸ As discussed in n. 6, this result is itself a consequence of some of the special assumptions that are typically imposed in these types of models, including ours, and can be relaxed. Since this is not our main focus, we do not explore this issue further in this paper.

¹⁹ Strictly speaking, proposition 8 requires the worker type to have exactly the same productivity profile as capital. In proposition 10, we extend this to worker types whose productivity profile is sufficiently similar—but not exactly equal—to that of capital. This extension is easier to formalize when studying global changes in capital productivity, so we defer it to proposition 10.

C. Productivity

The next proposition provides a characterization of the productivity effects of automation, using total factor productivity (TFP) as a measure for productivity. TFP growth in our setting is

$$\begin{aligned}\Delta \log \text{TFP} &= \Delta \log Y - \alpha_k \Delta \log \left(\frac{\bar{K}}{q} \right) - (1 - \alpha_k) \Delta \log \left(\int_0^1 l_s ds \right) \\ &= \Delta \log Y - \alpha_k \Delta \log \left(\frac{\bar{K}}{q} \right),\end{aligned}$$

where we subtract the growth in the value of the capital stock, $\Delta \log(\bar{K}/q)$, not the change in the quantity of machines, $\Delta \log \bar{K}$, from gross output growth. Greenwood, Hercowitz, and Krusell (1997) argue that one should subtract the growth of the quantity of capital if the goal is to obtain a measure of factor-neutral productivity growth. However, to obtain a measure of overall productivity change, including capital-specific productivity, one needs to subtract the growth of the value of capital, as they also point out. We opt for the latter since we want our productivity measure to reflect the capital-specific productivity gains due to higher q .

With this definition of TFP, we obtain the following second-order Taylor expansion.

PROPOSITION 9 (Productivity effects). Suppose that assumptions 1–3 and condition 1 hold, let $q \in (q_0, q_m)$, and consider a small increase in the productivity of capital $\Delta \log q > 0$. Then we have

$$\begin{aligned}\Delta \log \text{TFP} &\approx \alpha_k \Delta \log q \\ &+ \frac{\alpha_k}{1 - \alpha_k} \left[\lambda - 1 + \frac{\partial \log \Gamma_k}{\partial \bar{x}} \frac{d \bar{x}}{d \log q} + \frac{\partial \log \Gamma_k}{\partial \underline{x}} \frac{d \underline{x}}{d \log q} \right] (\Delta \log q)^2.\end{aligned}$$

The first term in the approximation is an immediate consequence of Hulten's theorem and implies that the first-order effect of an increase in capital productivity is equal to its income share. This first term indicates that the TFP gain will be smaller when α_k is small, thus confirming the result in proposition 8.2: when the productivity of capital is low to start with or, equivalently, only a few tasks are initially automated, then the productivity effect is small (which is the reason why negative wage effects are more likely in this case).

The second term captures two distinct but related forces. First, the expansion of the task share of capital Γ_k tends to make the TFP effects of lower capital costs (higher capital productivity) convex: a lower cost for capital expands the set of automated tasks, generating a bigger base on which additional productivity gains can be obtained. Second, holding the set of

automated tasks fixed, lower capital costs increase or decrease the share of capital in national income, depending on whether tasks are complements ($\lambda < 1$) or substitutes ($\lambda > 1$). If tasks are complements, a lower cost of capital leads to an increase in the labor share and a decrease in the capital share. This partially counters the effect from the expansion of the set of automated tasks and the implied convexity of TFP effects. In contrast, if tasks are substitutes, a lower cost of capital reduces the labor share and raises the capital share, thus amplifying convexity.

V. Global Effects of Automation

In this section, we consider noninfinitesimal (potentially large) changes in the productivity of capital. We distinguish between two cases. In the first (sec. V.A), after this change, automation still remains interior. In the second (sec. V.B), we transition from interior to low-skill automation. Finally, we also discuss additional comparative statics with respect to labor supply changes.

A. Nonlocal Changes with Interior Automation

PROPOSITION 10 (Polarization with large changes in automation). Suppose that assumptions 1–3 and condition 1 hold, let $q \in (q_0, q_m)$, and consider a potentially large increase in the productivity of capital $\Delta \log q > 0$, which still satisfies $\log q + \Delta \log q < \log q_m$. Let $\tilde{s}' = \tilde{s} + \Delta \tilde{s} \in (0, 1)$ be the new threshold skill level.

1. Then, automation expands in both directions and causes employment polarization. That is,

$$\Delta \underline{x} < 0 \text{ and } \Delta \bar{x} > 0,$$

and

$$\Delta x_s < 0 \text{ for all } s \in (0, \tilde{s}') \text{ and } \Delta x_s > 0 \text{ for all } s \in (\tilde{s}', 1).$$

2. There is wage polarization in the sense that skill premia increase above the threshold skill \tilde{s}' and decrease below this threshold. Or equivalently,

$$\begin{aligned} \Delta \log w_s &> \Delta \log w_{s'} \text{ for all } s < s' \in (0, \tilde{s}'], \\ \Delta \log w_s &< \Delta \log w_{s'} \text{ for all } s' > s \in [\tilde{s}', 1). \end{aligned}$$

3. The average wage always increases, and moreover, there exists a threshold for capital productivity $\hat{q} > q_0$ such that if $q + \Delta q \in (q_0, \hat{q})$, then for some $\delta_1, \delta_2 > 0$, we have $d \log w_s < 0$ for all $s \in [\tilde{s} - \delta_1, \tilde{s} + \delta_2]$.

4. Let $\gamma_{s,k}^{\max} = \max_x \{\log \psi_{s,x} - \log \psi_{k,x}\} - \min_x \{\log \psi_{s,x} - \log \psi_{k,x}\}$ for some s . Then,

$$\Delta \log w_s \leq \gamma_{s,k}^{\max} - \Delta \log q.$$

In particular, if $\gamma_{s,k}^{\max} < \epsilon$, then any $\Delta \log q > \epsilon$ will reduce the wage of workers with skill level s .

5. Suppose that $\psi_{0,0}/\psi_{k,0} < \psi_{0,1}/\psi_{k,1}$. Then, if $q + \Delta q$ is sufficiently close to q_m , the inequality between the top and the bottom of the skill space increases, that is,

$$\Delta \log w_0 < \Delta \log w_1.$$

In summary, this proposition establishes that our main employment and wage polarization results do not depend on whether we consider small or large changes in capital productivity, provided that automation starts out and remains interior. Moreover, as before, when we initially have relatively few tasks automated (or the productivity of capital is still relatively low), an expansion in automation hurts workers around the skill threshold \tilde{s} . Our refinement of Wiener's conjecture also extends to this case: the wages of worker types with productivity profiles sufficiently similar to capital's will decline (but, as before, wages cannot decline for all worker types). Finally, under the same conditions as in the local analysis, the impact of automation on the wage distribution is asymmetric: inequality increases between high-skill and low-skill workers. The proof of this proposition and of the remaining results are in appendix B.

We will next see that when automation ceases to be interior, we obtain very different comparative statics.

B. Transition to Low-Skill Automation

We next consider a nonlocal change in capital productivity inducing a transition from interior to low-skill automation.

PROPOSITION 11 (Transition to low-skill automation). Suppose that assumptions 1–3 and condition 1 hold. Suppose also that condition 2 holds and that $q \in (q_0, q_m)$ initially. Now consider a potentially large increase in the productivity of capital $\Delta \log q > 0$ such that $\log q + \Delta \log q > \log q_m$. Then:

1. Automation transitions from interior to low skill, so all low-complexity tasks are taken over by capital. That is, $\Delta \underline{x} = -\underline{x}$.
2. This transition does not induce employment polarization. Instead, the assignment function shifts up everywhere, $\Delta X_s > 0$ for all $s < 1$.

3. The transition does not induce wage polarization either. Instead, skill premia increase over the entire skill space. That is,

$$\Delta \log w_s < \Delta \log w_{s'} \text{ for all } s < s'.$$

Next, suppose that assumptions 1–3 and conditions 1 and 2 hold but we start from $q \geq q_m$ already. Then, a further increase in the productivity of capital shifts up the assignment function everywhere, and skill premia increase over the entire skill space (i.e., there is no longer any employment or wage polarization).

The most important result in this proposition is that for a sufficiently large capital productivity, automation transitions from interior to low skill, and this transition changes the wage effects of automation qualitatively. In contrast to the wage polarization pattern we have seen so far, once automation becomes low skill, further automation induces monotone increases in wage inequality, impacting the lowest-skill workers most negatively (or least positively). In this case, the employment polarization effects of automation vanish as well: further automation now pushes all workers toward more complex tasks. Figure 3 diagrammatically illustrates this transition.

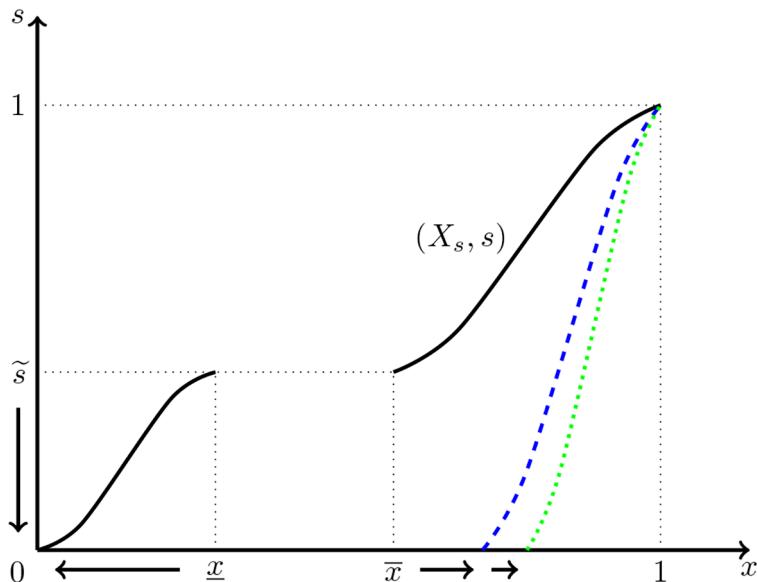


FIG. 3.—Transition to low-skill automation. If capital productivity grows sufficiently, automation becomes low skill (dashed line). A further increase in capital productivity then pushes all workers toward the upper end of the task distribution (dotted line).

C. Implications of Labor Supply Changes

Finally, we consider the wage implications of changes in the labor supply profile. Although this could have been studied with local analysis, it is more convenient to discuss these comparative statics in the case of global changes. The main comparative static is given in the next proposition.

PROPOSITION 12 (Labor supply changes and automation). Suppose that assumptions 1–3 as well as conditions 1 and 2 hold and let $q \in (q_0, q_m)$ under the initial labor supply l . Now consider a change in labor supply such that $\Delta \log l_s < \Delta \log l_{s'}$ for all $s < s'$ (an increase in the relative supply of more skilled workers), and suppose that the resulting decline in q_m ($\Delta q_m < 0$) leads to $q_m + \Delta q_m < q$ such that automation transitions from interior to low skill. Then, skill premia increase in the bottom part of the wage distribution and decrease in the upper part. Specifically, there exists $\hat{s} \in (\tilde{s}, 1)$ (where \tilde{s} is the threshold skill before the labor supply change) such that

$$\begin{aligned}\Delta \log w_s &< \Delta \log w_{s'} \text{ for all } s < s' \in (0, \hat{s}], \\ \Delta \log w_s &\geq \Delta \log w_{s'} \text{ for all } s' > s \in [\hat{s}, 1).\end{aligned}$$

Moreover, $\Delta \log w_0 > 0$, meaning that the real wage at the bottom of the distribution increases.

The proposition complements our analysis of labor supply changes in proposition 5. There, we showed that an increase in relative skill supply can trigger a transition from interior to low-skill automation by raising low-skill workers' wages and making it more profitable to automate their jobs. In this case, proposition 12 establishes a type of upward-sloping relative demand for skills. Given a fixed assignment of workers and capital to tasks, the increase in the supply of skills would have reduced skill premia. However, the response of equilibrium assignment alters this pattern qualitatively. Low-skill automation becomes more likely, and this reduces the relative wages of low-skill workers and raises skill premia at the bottom. Other instances of greater relative supply of skills leading to higher skill premia are present in models of directed technological change (because greater abundance of skilled workers encourages more skill-biased technological change, as in Acemoglu [1998, 2007]) and in models of search and matching (because with more skilled workers around, more employers make investments complementary to skilled workers and search for them, as in Acemoglu [1999]). In the model here, a similar outcome arises, even though there is no endogenous innovation and all markets are competitive. Rather, this result is driven by the response of the equilibrium assignment of tasks between capital and labor.

Notice also that the proposition establishes $\Delta \log w_0 > 0$, which means that low-skill wages never fall in absolute terms in response to an increase

in the relative supply of more skilled workers. This implies that although the endogenous transition from interior to low-skill automation has a negative effect on low-skill wages, this effect is not strong enough to offset the positive direct impact coming from the increased relative supply of more skilled workers, and consequently, the demand for the low-skill workers is always downward sloping.

VI. Quantitative Analysis

In this section, we undertake a preliminary quantitative analysis of the consequences of different types of automation and policies in our framework.

A. Calibration

We normalize labor supply to one for all skill types, that is, $l_s \equiv 1$, which implies that the skill index s represents percentiles of the wage distribution. For labor productivity, $\psi_{s,x}$, we start from the specification in Teulings (1995, 2005), where productivity is assumed to be log linear in the product of s and x . We extend this by including a quadratic in s :

$$\log \psi_{s,x} = ms + ns^2 + asx + A,$$

where $a > 0$ corresponds to more skilled workers having comparative advantage in more complex tasks. The quadratic in s regulates the extent of the absolute advantage of more skilled workers, which is useful for matching the empirical wage distribution. In particular, without the ns^2 term, this functional form generates too much inequality at the top of the wage distribution because of the linear absolute advantage of more skilled workers.

The simplest specification for capital productivity, $\psi_{k,x}$, consistent with assumption 3, is

$$\log \psi_{k,x} = a_k x + bx^2,$$

with $b < 0$.

Given these specifications, we have to choose eight parameters: m , n , a , a_k , b , the capital productivity parameter q , the elasticity of substitution between tasks λ , and the labor-augmenting technology parameter A . We set $\lambda = 0.5$ externally, using the estimate of the firm-level elasticity of substitution between production workers, tech workers, and other workers from Humlum (2021), which we interpret (following Humlum) as the firm-level elasticity of substitution between tasks. Next, we choose m , n , a , a_k , b , and q to match two sets of empirical moments. The first are moments from the US income distribution in 1980, including the 90–50 and the 50–10 differences in log hourly wage, and the factor income share

of equipment and software capital in 1980. We focus on the share of equipment and software capital because all capital is used for automation in our model. The second set of moments concerns the impact of automation on the US wage distribution between 1980 and 2016–17, which is from the estimates in Acemoglu and Restrepo (2022). Specifically, we take from that paper the estimates of the impact of automation on the 30–10, 50–30, and 90–50 differences in log hourly wages. We include the 30th percentile because this is where automation had its least positive (most negative) impact on wages according to their estimates (see fig. 5). Additionally, we target the change in the income share of equipment capital and software between 1980 and 2016. We reduce the cost of capital $1/q$ in the model by 79%, which corresponds to the decline in the real cost of equipment and software in the United States between 1980 and 2016. We choose the labor-augmenting technology parameter A to match the level of wages in 1980 (this parameter does not affect any of the relative equilibrium quantities).

We describe the procedure for numerically solving for the equilibrium of our model in appendix section B.5.1, our data sources and the construction of empirical moments in appendix section B.5.2, and the details of the calibration procedure in appendix section B.5.3.

The results are displayed in table 1. Panel A shows the calibrated parameter values, while panel B depicts the aforementioned moments both in the data and in our calibration. The model fits all data moments closely. Figure 4 additionally shows that the model matches the 1980 US wage distribution very well even beyond the targeted percentiles. In figure 5, we present the log wage changes induced by automation between 1980 and 2016 in our model and in the data (fig. 5A) as well as the corresponding change in the assignment function (fig. 5B). As expected, the set of automated tasks expands in both directions, while the wage distribution becomes more polarized. The wage effects predicted by our model track the changes in the data quite well, but the two lines differ by a constant of about 0.3, on average: the model predicts a median wage increase of 0.3 log points, while the estimated series shows a median wage change close to zero. This discrepancy is likely due to the fact that the estimates in Acemoglu and Restrepo (2022) isolate the effects of expansions in the set of automated tasks, while an increase in q in our model both expands automation and induces deepening of automation, further raising productivity in already automated tasks, thus amplifying the productivity effect and pushing in the direction of positive wage growth for all skill groups. Consistent with this, the implied median wage increase of 0.3 log points in our model is in the ballpark of the actual change in US median wages (of about 0.26 log points), though the actual median wage change is likely impacted by a variety of other factors, including other types of technological changes, educational upgrading, new tasks, and institutional and

TABLE 1
 CALIBRATION PARAMETERS AND RESULTS

Parameter/Moment	(1)	(2)
A. Calibrated Parameters		
	Value	Rationale/Target
λ	.5	Humlum 2021
A	-.20	log median wage
m	2.6	Jointly calibrated
n	-2.5	Jointly calibrated
a	8.0	Jointly calibrated
a_k	100	Jointly calibrated
b	-310	Jointly calibrated
$\log q$	-9.3	Jointly calibrated
B. Comparison of Moments		
	Data	Model
log 50–10 wage percentile ratio, $\log(w_{.5}/w_{.1})$.60	.60
log 90–50 wage percentile ratio, $\log(w_{.9}/w_{.5})$.67	.67
Income share of equipment and software, α_k	.16	.17
Change in log 30–10 wage percentile ratio due to automation, $\Delta \log(w_{.3}/w_{.1})$	-.02	-.01
Change in log 50–30 wage percentile ratio due to automation, $\Delta \log(w_{.5}/w_{.3})$.03	.03
Change in log 90–50 wage percentile ratio due to automation, $\Delta \log(w_{.9}/w_{.5})$.16	.16
Change in income share of equipment and software from 1980 to 2016, $\Delta \alpha_k$.01	.03
Log median wage level (in 2008 dollars), $\log w_{.5}$	2.6	2.6

NOTE.—Panel A shows the calibrated parameter values (col. 1) and the targets used in their calibration (col. 2). Jointly calibrated parameters are set to match all moments except for the log median wage level. Panel B presents the data moments used in the calibration (col. 1) and their model counterparts (col. 2). See app. sec. B.5.2 for data sources.

norm changes in the US labor markets. In this context, we also note that our calibration does not impose conditions 1 and 2, but both of these conditions are comfortably satisfied given the implied parameter values.

Although there is no one-to-one mapping between the seven moments we use from the United States and the six parameters, each parameter is closely related to a particular equilibrium outcome of the model, which allows us to gain additional intuition about the parameter choices and the consequences of reasonable variations therein. The parameters m and n determine the degree of absolute advantage among workers. A higher m

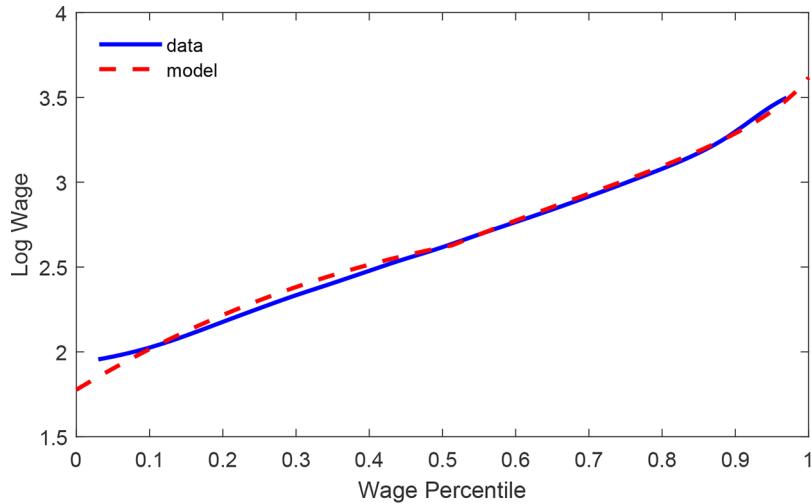


FIG. 4.—Log wages by percentile of wage distribution in 1980. The solid line corresponds to the data (based on Acemoglu and Autor 2011), while the dashed line is computed from our calibration, using the parameter values shown in table 1. For data sources, see appendix section B.5.2.

raises skill premia across the entire wage distribution, while increasing n does so mainly at the upper end. Choosing m and n jointly thus allows us to match both the 90–50 and 50–10 wage ratios in 1980.

The parameter a determines the degree of comparative advantage among workers and their substitutability. A high value of a corresponds to low substitutability between skill types and implies that the displacement effect of automation is largely concentrated on those workers who are directly displaced. Graphically, this corresponds to the two branches of the dashed line in figure 5A becoming steeper. Thus, by varying a , we can scale the effects of automation on the 30–10, 50–10, and 90–50 wage ratios up or down.

The parameters a_k and b determine the slope and curvature of the capital productivity curve. The slope parameter controls where in the task space capital has its greatest comparative advantage relative to workers and thus where the set of automated tasks is located. A higher slope moves the set of automated tasks toward the top of the wage distribution. Graphically, this means that a_k regulates where the minimum of the dashed line is in figure 5A.

The curvature parameter b , on the other hand, determines how quickly the productivity of capital declines as we move away from the set of automated tasks and, as such, controls by how much the set of automated tasks expands when capital costs fall further. The value of b is therefore disciplined directly by the increase in the income share of equipment

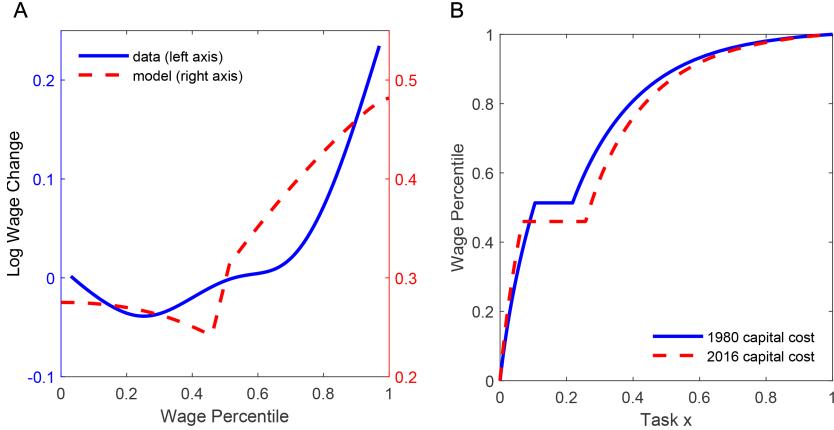


FIG. 5.—Change in log wages and assignment due to automation, 1980–2016. A, Estimated change in log wages between 1980 and 2016 due to automation from Acemoglu and Restrepo (2022; solid line) and corresponding changes from our calibrated model (dashed line). B, Corresponding change in assignment function in our model. The values of parameters used in the calibration are shown in table 1. For data sources, see appendix section B.5.2.

and software capital in the data. The level of the capital cost $1/q$ determines the extent of automation in the baseline and is thus set to match the 1980 equipment and software income share in the initial equilibrium.

Finally, as noted above, because the parameter A scales only wages up or down, it is not calibrated jointly with the rest of the parameters and is set at the end of the calibration to match the median wage in 1980 exactly given the other parameters.

B. Robustness

We check the robustness of our calibration in two dimensions. First, in our baseline, the elasticity of substitution between tasks, λ , is set to 0.5 following the firm-level elasticity of substitution estimate of Humlum (2021). Even though we believe that firm-level task substitution is a good approximation to the elasticity of substitution between tasks at the aggregate, this aggregate elasticity could also differ from the firm-level elasticities because of adjustment costs or heterogeneities at the firm level. Our first exercise is to verify that reasonable variations in this elasticity do not majorly impact our results. Specifically, we change λ to 0.3 or 0.7, and the implications for targeted moments are depicted in table 2. Table 2 and figure 6A show that the impact on the implied (log) wage changes due to automation between 1980 and 2016 remains fairly similar, and therefore our calibrated model continues to match the data reasonably well.

TABLE 2
ROBUSTNESS CHECKS

Moment	Data	Baseline	Elasticity of Substitution		Capital Price Index	
			$\lambda = .3$	$\lambda = .7$	Equipment	Software
$\log(w_5/w_1)$.60	.60	.57	.70	.60	.60
$\log(w_9/w_5)$.67	.67	.46	.85	.67	.67
α_k	.16	.17	.19	.15	.17	.17
$\Delta \log(w_3/w_1)$	−.02	−.01	−.01	−.01	−.01	−.01
$\Delta \log(w_5/w_3)$.03	.03	−.03	.08	.05	.02
$\Delta \log(w_9/w_5)$.16	.16	.19	.13	.19	.15
$\Delta \alpha_k$.01	.03	.01	.05	.02	.03
$\log w_5$	2.6	2.6	2.5	2.8	2.6	2.6

NOTE.—The table shows the targeted moments in the data under our baseline calibration for alternative values of the elasticity of substitution across tasks and for alternative price indexes used to compute the decline of capital costs between 1980 and 2016.

Second, we vary the capital cost decline. In the baseline, we chose a cost decline of 79%, building on the quality-adjusted price index for nonresidential equipment and software provided by DiCecio (2009) and used in many other studies, such as Grossman et al. (2017) and Hubmer (2023). In table 2 and figure 6B, we report how our calibration targets are affected when we use two alternative price indexes to construct our capital cost decline. First, we use the quality-adjusted price index by DiCecio (2009) for just nonresidential equipment, which yields a capital cost decline of 85%. Second, we use the price index for software from the Bureau of Economic Analysis (which does not include a full quality adjustment). This index yields a cost decline of 74% (see app. sec. B.5.2 for details on how we compute these cost declines). Table 2 and figure 6B show that our model matches the data well with these alternative price indexes.²⁰

C. Counterfactuals

We now use our calibrated model for four counterfactual exercises that shed light on how further automation, different types of technological changes, minimum wages, and changes in labor supply might impact the wage distribution.

We first consider the consequences of reducing capital costs: increasing capital productivity. To make comparisons easier, we focus on a further 79% decline in the cost of capital, $1/q$ (equivalent to the decline between

²⁰ When we vary the elasticity of substitution between tasks or the capital cost decline, we recompute our targeted moments in table 2 and the log wage changes in fig. 6, holding all internally calibrated parameters at their levels from our baseline calibration. If we recalibrated the internal parameters for any given alternative value of the external parameters, the fit in table 2 and fig. 6 would be even better. We do not undertake this recalibration in order to show more clearly the implications of varying the external parameters (while holding the internal parameters fixed).

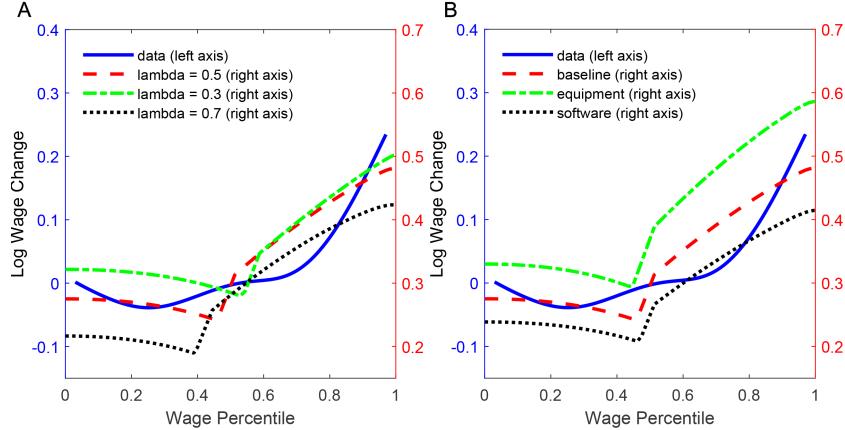


FIG. 6.—Model-implied changes in log wages due to automation between 1980 and 2016 for different values of elasticity of substitution λ (A) and different price indexes used to compute decline in cost of capital from 1980 to 2016 (B). All other parameters are fixed at their values from our baseline calibration (see table 1).

1980 and 2016 in the United States). Unsurprisingly, the results (fig. 7) are similar to the analogous wage changes between 1980 and 2016–17. There is again a sizable amount of wage polarization.²¹ Nevertheless, there are some important differences as well. Automation moves further toward the bottom of the wage distribution than it did between 1980 and 2016, and hence displacement falls more on the shoulders of low-skill workers, boosting inequality even more than in 1980–2016. This result accords with proposition 10, which established that further declines in the cost of capital that take us closer to the threshold for low-skill automation will tend to increase inequality between the top and the bottom of the wage distribution. Put differently, given the global comparative advantage of capital in condition 2, further declines in the cost of capital leave lower-skill workers more vulnerable to automation.

There are in principle many ways in which automation opportunities could improve. This raises the possibility that the next wave of automation may have very different consequences than what we experienced between 1980 and 2016. The most important reason for this may be advances in (generative) AI, which could expand the reach of automation to a broader

²¹ As explained in sec. VI.B, reducing the cost of capital generates a significant amount of deepening of automation in our model, which is the reason why in these counterfactuals wages increase even at the bottom of the distribution. An alternative, in line with our discussion above, would be to remove these nonautomation implications by reducing A by an equivalent amount. In that case, the median equilibrium wage would remain roughly constant following the decline in the cost of capital, and there would be sizable real wage declines at the bottom of the distribution.

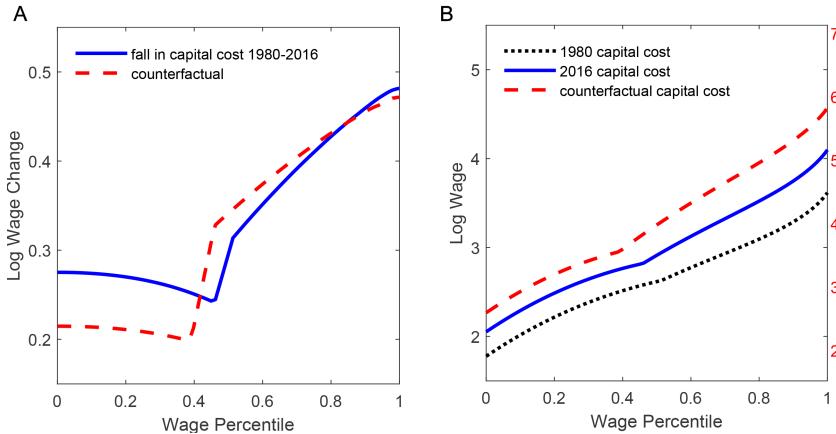


FIG. 7.—Change in wages due to counterfactual fall in capital costs. A, Change in log wages in response to further fall in capital cost by 79% starting from its 2016 level (dashed line) and, for comparison, our baseline decline in cost of capital (solid line). B, Corresponding log wages (in level) by percentile of wage distribution (dotted line = 1980 baseline; solid line = 2016 capital cost; dashed line = counterfactual). The values of parameters used in the calibration are shown in table 1. For data sources, see appendix section B.5.2.

set of occupations. Generative AI may have less impact on already automated tasks—such as software-based clerical functions and various production tasks, such as assembly, welding, and painting—that have been automated by robots and advanced automatic equipment. This motivates us to consider an AI counterfactual where the curvature of the capital productivity schedule $\log \psi_{k,x}$ is reduced but its maximum remains unaffected, so that capital becomes more productive at the tails of the task distribution but not so much in already automated tasks. Formally, we benchmark capital productivity against the productivity of the median worker, $\log \psi_{k,x} - \log \psi_{0.5,x}$, and reduce the curvature of this quadratic function while keeping its maximum fixed by adjusting parameters a_k and b (see app. sec. B.5.4 for details). We further discipline the exact shape of the new productivity schedule by imposing that the aggregate net output gains—and thus the productivity effects—must be the same as from a uniform capital cost reduction of 79%, as in our first counterfactual.²²

The wage consequences of the AI calibration are shown in figure 8. As a benchmark, figure 8 also includes the results from our first counterfactual

²² There is considerable uncertainty about which tasks AI will impact and how much it will increase productivity. Acemoglu (2024) estimates that less than 5% of the economy will experience automation or significant productivity increases due to AI within the next 10 years, while others estimate much more extensive effects of AI. Acemoglu (2024), like Eloundou et al. (2024), also argues that AI will mostly impact previously nonautomated cognitive tasks, which is consistent with our assumption here.

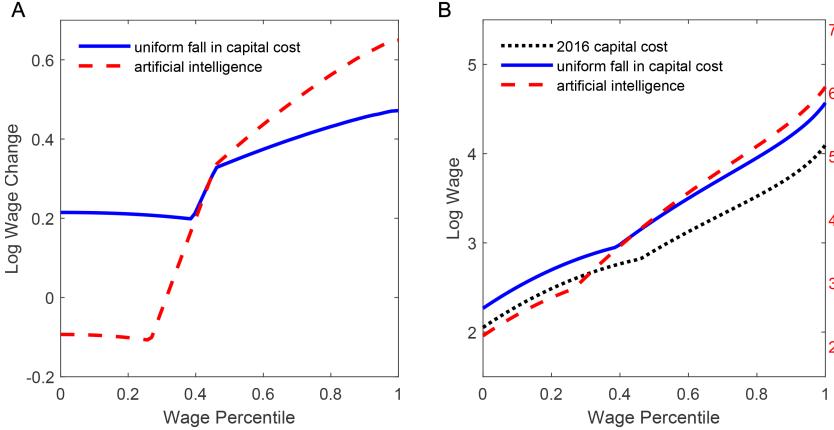


FIG. 8.—Change in wages due to counterfactual progress in AI. *A*, Change in log wages in response to change in capital productivity mimicking effects of advances in generative AI (dashed line) and, for comparison, uniform capital productivity increase counterfactual from figure 7 (solid line). *B*, Corresponding log wage levels by wage percentile. The values of parameters used in the calibration are shown in table 1.

where capital productivity increases uniformly. While both counterfactuals result in employment and wage polarization, their wage effects are quite distinct. AI leads to steeper wage declines at the bottom of the wage distribution (because automation expands even further at the bottom) and bigger wage gains at the top (since higher-skilled workers who do not suffer automation benefit from the productivity gains). This is a consequence of the way we have modeled the AI advances: less deepening of automation and productivity gains in already automated tasks and more aggressive expansions in previously nonautomated tasks.

Finally, we explore the implications of an increase in the minimum wage and a reduction of low-skill labor supply in this setup. For both counterfactuals, we fix the capital cost at its 2016 level. We start with the effects of an increase in the minimum wage. We take the employment-weighted average minimum wage level across US states in 2022, which was \$10.35, and increase this to the maximum of state-level minimum wages in 2022 (\$16.1 in Washington, DC). When computing the employment-weighted average minimum wage across states, we take into account that the federal minimum wage is binding in states with a minimum wage below the federal level (see app. sec. B.5.2 for details). The higher minimum wage leads to the automation of low-skill tasks as predicted by proposition 4 (see fig. B-4 [figs. B-1 through B-5 are available online]). The wage implications of the minimum wage are displayed in figure 9, which shows a significant impact on the entire wage distribution, with wage increases of up to 10% at

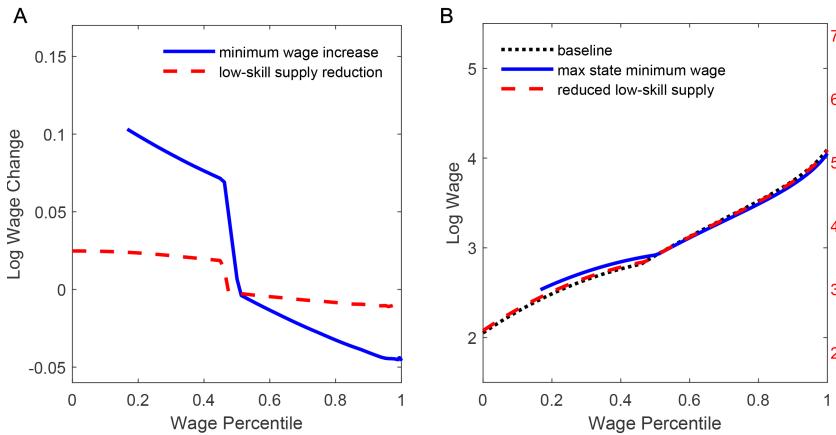


FIG. 9.—Change in wages due to minimum wage increase and low-skill supply reduction. *A*, Change in log wages when minimum wage rises from its cross-state average to maximum state level in 2022 (\$16.1; solid line) and change in log wages following reduction in labor supply of low-skilled workers (dashed line). *B*, Corresponding log wage levels by (initial) wage percentile. The values of parameters used in the calibration are given in table 1, and capital productivity is kept at its 2016 level.

the bottom and declines of up to 5% at the top.²³ The negative impact at the top comes from the fact that the minimum wage leads to an approximately 10% reduction in employment, entirely concentrated at the bottom. This reduces aggregate output and lowers demand for labor across all tasks.

Next, we consider a reduction in the supply of low-skill workers, which might result from drastic restrictions of immigration from low-income countries. We reduce the supply of workers in the bottom 20 percentiles of the wage distribution uniformly by 10% and phase this reduction out linearly between the 20th and the 40th percentiles of the wage distribution, leading to an overall decline in labor supply of 3% (see app. sec. B.5.5 for details). Consistent with proposition 5 and the empirical results in Clemens, Lewis, and Postel (2018), this change also induces low-skill automation (see fig. B-4). The wage effects, also presented in figure 9, show a more modest impact from this labor supply change compared with the minimum wage increase: wages for low-skill workers increase by about 3%, while wages at the top of the distribution fall by slightly more than 1%. Overall, the minimum

²³ In figs. 4–8, the horizontal axis is the current wage percentile (which is also identical to the initial wage percentile of the workers) since throughout, wages are monotone in skills, and in these experiments, there is always full employment and there are no new workers. When we increase the minimum wage or consider changes in labor supply, initial and actual wage percentiles no longer overlap, and in fig. 9, we use the initial wage percentile on the horizontal axis, which facilitates interpretation.

wage hike appears quite powerful in reducing wage inequality in our setting, but it also leads to a decline in total employment by 10% compared with a reduction of only 3% under our counterfactual low-skill supply.

VII. Conclusion

There has been rapid automation of a range of tasks across the industrialized economies over the past four decades. There is growing evidence that this automation has fueled both inequality and polarization, with middle-skilled workers being displaced from their jobs and experiencing relative (and sometimes absolute) wage declines.

To develop a deeper understanding of the causes of polarization, this paper has built an assignment model of automation. In our model, each of a continuum of tasks of variable complexity is assigned to either capital or one of a continuum of labor skills. Our model generalizes existing assignment models, which typically impose global supermodularity conditions that ensure monotone matching between factors and tasks. In contrast, in our model with capital, there is no global supermodularity.

We prove existence and uniqueness of competitive equilibria and characterize conditions under which automation is interior, meaning that it is tasks of intermediate complexity that are assigned to capital. In a nutshell, interior automation arises when the most skilled workers have a comparative advantage in the most complex tasks relative to capital and other labor and when the wages of the least skilled workers are sufficiently low relative to their productivity and the effective cost of capital in low-complexity tasks, so that it is not profitable to use capital or algorithms instead of low-skill workers. Highlighting the role of wages at the low-end of the wage distribution, we demonstrate that minimum wages and other sources of higher wages at the bottom make interior automation less likely.

We provide a series of local and global comparative statics, showing how further automation impacts wages and assignment patterns. Most importantly, when automation starts and remains interior, a lower cost of capital (or greater capital productivity) causes employment polarization: middle-skill workers are displaced from middle-complexity tasks and are pushed toward higher or lower parts of the complexity distribution. This type of automation also causes wage polarization: the skill premium monotonically increases above a skill threshold and monotonically declines below the same threshold. Moreover, automation tends to reduce the real wage of workers with comparative advantage profiles close to that of capital.

Our global comparative static results additionally establish that a large enough increase in capital productivity ultimately induces a transition to low-skill automation, whereby the pattern of comparative statics changes qualitatively. In particular, after this transition to low-skill automation,

further declines in the cost of capital no longer cause employment or wage polarization. Rather, they have a monotone effect on the skill premium.

Despite its richness, our framework is tractable and opens the way to further analyses of the changing assignment patterns in modern labor markets. We illustrated this richness by presenting both a range of comparative static results and a simple calibration exercise where we explored the quantitative implications of various counterfactual technological and institutional changes.

There are many areas for future fruitful inquiry. First, automation has been going on together with a changing structure of tasks and an evolving distribution of skills over at least the past 250 years. This can be introduced into our framework by simultaneously expanding the range of tasks and skills and would be an important area for new work. Second, the productivity of capital in various tasks should in principle change endogenously, responding to which tasks are being assigned to capital or are likely to be assigned to capital in the future. This issue can be investigated in an extended version of our framework, in which the direction of technological change and capital productivity across tasks are endogenized. Third, in practice, multiple tasks may be assigned to a worker because there are either economies of scope or other types of task complementarities, and extending this class of models to one-to-many matching is another important area for further inquiry. Fourth, our quantitative exploration can be significantly expanded, for example, by introducing a more general family of comparative advantage schedules and then estimating the parameters of this family using more moments from the baseline wage distribution and the response of the wage distribution to different types of shocks. There are also several other counterfactual exercises to be considered, including those related to changes in supplies and offshoring-type opportunities. Last but not least, the framework here can be used to further refine the empirical investigation of the relationship between automation and inequality, for example, by adding more structure and predictions to studies such as Acemoglu and Restrepo (2022).

Appendix A

A1. A Converse to Proposition 3

In proposition 3, we show that automation is interior for low levels of capital productivity if condition 1 holds (together with assumptions 1–3). Here we provide an inverse to this result: if condition 1 is violated, automation is either high or low skill (again under assumptions 1–3). In this sense, condition 1 is necessary for interior automation in our setup.

PROPOSITION A1 (Interior automation, converse). Suppose that assumptions 1–3 hold. Then, we have the following:

1. If condition 1.1 is violated while condition 1.2 holds, then automation is low skill for all $q \in (q_0, q_2)$.
2. If condition 1.1 holds while condition 1.2 is violated, then automation is high skill for all $q \in (q_0, q_2)$.

Proof. We focus on the case where condition 1.1 is violated while condition 1.2 holds, as the proof for the second case is symmetric.

If condition 1.1 does not hold, we have

$$\frac{\partial \log \psi_{k,0}}{\partial x} - \frac{\partial \log \psi_{0,0}}{\partial x} \leq 0.$$

By log supermodularity of labor productivity (assumption 2), it follows that

$$\frac{\partial \log \psi_{k,0}}{\partial x} - \frac{\partial \log \psi_{s,0}}{\partial x} < 0 \text{ for all } s > 0.$$

Together with quasi-concavity of $\psi_{k,x}/\psi_{s,x}$ (assumption 3), this implies that $w_s \psi_{k,x}/\psi_{s,x}$ is decreasing in x for all $s > 0$. By continuity of productivity schedules, the same must then hold for $s = 0$. Hence, the minimal effective unit cost of producing the amount $\psi_{k,x}$ of task x with labor, $\omega_x = \min_s \{w_s \psi_{k,x}/\psi_{s,x}\}$, is the lower envelope of a family of decreasing functions and must therefore be decreasing itself. This implies that $\{x | \omega_x \geq 1/q\}$ either is empty (if $q < q_0$) or contains zero (if $q \geq q_0$). Using lemma A1, the same holds for the set of automated tasks X_k . QED

A2. Proofs for Section III: Equilibrium and Interior Automation

A2.1. Proof of Proposition 1: Existence and Uniqueness

Existence.—An equilibrium allocation maximizes net output subject to labor market clearing, given by

$$\int_0^1 L_{s,x} dx \leq l_s \text{ for all } s. \quad (\text{A1})$$

To prove existence, it is useful to split the problem of net output maximization into two steps. First, we fix the aggregate capital stock \bar{K} and maximize gross output for given \bar{K} and subject to (A1). This is a problem of maximizing a continuous function over a compact set such that a maximizer is guaranteed to exist. Let $F(\bar{K}, l)$ denote the maximal gross output for given \bar{K} and labor supply l .

In the second step, we choose \bar{K} to maximize net output $F(\bar{K}, l) - \bar{K}/q$. This is again a continuous problem, but \bar{K} can be any positive real number, so we have to establish boundedness. For this, note that $\lim_{\bar{K} \rightarrow \infty} \partial F(\bar{K}, l) / \partial \bar{K} = (\int_0^1 \psi_{k,x}^{\lambda-1} dx)^{1/(\lambda-1)}$. Thus, assumption 1 ensures that $\lim_{\bar{K} \rightarrow \infty} \partial F(\bar{K}, l) / \partial \bar{K} < 1/q$ such that net output is bounded and attains its maximum for finite \bar{K} .

Essential uniqueness.—For essential uniqueness of equilibrium, note that net output is concave in the allocation and the set of feasible allocations is convex. This implies that while the equilibrium allocation itself may not be unique, the

Fréchet derivative of net output is constant across all equilibrium allocations.²⁴ Hence, equilibrium wages are unique. The same argument applies to task prices when writing the maximization of net output as a maximization over task inputs, including task production functions as constraints, that is,

$$\max_{\{Y_x\}_{x=0}^1, L, K} \left[\int_0^1 Y_x^{(\lambda-1)/\lambda} dx \right]^{\lambda/(\lambda-1)} - \frac{1}{q} \int_0^1 K_x dx$$

subject to task production (2) and labor market clearing (3).

Uniqueness.—Our proof of proposition 2 shows that under assumptions 1–3, the labor allocation L and the set of tasks performed by capital X_k are uniquely determined given wages. Moreover, given prices, X_k , and the labor allocation, choosing the output-maximizing capital allocation K is a strictly concave problem with a unique solution.

Thus, given wages and task prices, the equilibrium allocation is determined uniquely. Since equilibrium wages and task prices are unique, the equilibrium allocation is unique as well.

A2.2. Proof of Proposition 2: Convexity of Assignment

Monotonicity.—Monotonicity of the labor allocation under comparative advantage assumptions is a standard result. One way to prove it, which is useful for our argument in the next step, is presented here. Let

$$S_x^{\min} = \operatorname{argmin}_s \{\log w_s - \log \psi_{s,x}\}$$

be the set of skills that produce task x at minimal cost. By assumption 2, $\log w_s - \log \psi_{s,x}$ is strictly submodular, so Topkis's (Topkis 1998) monotonicity theorem implies that if $s \in S_x^{\min}$, $s' \in S_x^{\min}$, and $x > x'$, then $s \geq s'$. Moreover, if for some x there exist $s, s' \in S_x^{\min}$ with $s > s'$, then all skill levels in (s', s) can be assigned to only x . This creates a mass point in the density of labor over tasks such that $p_x = 0$, contradicting condition (4).²⁵ Hence, S_x^{\min} is a singleton for all x . Inverting this correspondence, we obtain $\tilde{X}_s = \{x | s \in S_x^{\min}\}$, which is a superset of X_s , $X_s \subseteq \tilde{X}_s$ for all s . From the properties of S_x^{\min} , it follows immediately that if $x \in \tilde{X}_s$, $x' \in \tilde{X}_{s'}$, and $s > s'$, then $x > x'$. Finally, since $X_s \subseteq \tilde{X}_s$ for all s , the same implication holds for X_s .

Convexity.—We start with the following lemma, which will be useful to establish properties of X_k throughout the paper.

LEMMA A1. Suppose that assumptions 1–3 hold and let

$$\omega_x = \min_s \left\{ \frac{w_s \psi_{k,x}}{\psi_{s,x}} \right\}$$

²⁴ The set of maximizers of a concave function on a convex set is a face of the hypograph of the function. Thus, there exists a supporting hyperplane of the hypograph that contains the entire set of maximizers. Together with differentiability, this immediately implies that the derivative of the function is constant on the set of maximizers.

²⁵ Note that this holds independent of whether capital is used in the production of task x . In any case, assigning a strictly positive measure of skills to x creates a mass point in task output at x , which leads to $p_x = 0$.

be the minimal effective unit cost of producing the amount $\psi_{k,x}$ of task x with labor. Then, the set of automated tasks is equal to the upper level set of ω at level $1/q$:

$$X_k = \left\{ x \mid \omega_x \geq \frac{1}{q} \right\}.$$

Proof. The unit cost of producing the amount $\psi_{k,x}$ of task x with capital is $1/q$. Therefore, we must have $1/q \leq \omega_x$ on X_k . Moreover, if $1/q < \omega_x$ at some x , then x must be in X_k . Hence,

$$\left\{ x \mid \omega_x > \frac{1}{q} \right\} \subseteq X_k \subseteq \left\{ x \mid \omega_x \geq \frac{1}{q} \right\}.$$

Now suppose that X_k and $\{x \mid \omega_x \geq 1/q\}$ differ by a set of strictly positive measure. Then, since the labor endowment has no mass points, a strictly positive measure of skills must be assigned to a subset of $\{x \mid \omega_x = 1/q\}$. In particular, there must exist skill levels $s_1 < s_2 < s_3$ assigned to tasks $x_1 < x_2 < x_3$ in $\{x \mid \omega_x = 1/q\}$. Moreover, since the cost-minimizing skill S_x^{\min} is unique for every task (see first step of the proof), we must have

$$\frac{w_{s_2}\psi_{k,x_1}}{\psi_{s_2,x_1}} > \frac{w_{s_1}\psi_{k,x_1}}{\psi_{s_1,x_1}} = \frac{w_{s_3}\psi_{k,x_2}}{\psi_{s_3,x_2}} = \frac{w_{s_2}\psi_{k,x_3}}{\psi_{s_2,x_3}} < \frac{w_{s_3}\psi_{k,x_3}}{\psi_{s_3,x_3}}.$$

But this string of relations contradicts the quasi-concavity of $\psi_{k,x}/\psi_{s,x}$. Hence, the difference between X_k and $\{x \mid \omega_x \geq 1/q\}$ must be of measure zero, and we can set $X_k = \{x \mid \omega_x \geq 1/q\}$ without loss of generality. QED

The “if” part of proposition 2.2 follows from lemma A1. By assumption 3, $w_s\psi_{k,x}/\psi_{s,x}$ is quasi-concave in x for all s . Thus, ω_x is the lower envelope of quasi-concave functions, and as such, it is quasi-concave itself. Hence, its upper level sets are convex and so is X_k .

Next, consider the “only if” part. We will prove that if $\psi_{k,x}/\psi_{s,x}$ is not quasi-concave in x for some s , then there exists a labor endowment l and capital productivity q such that X_k is not convex. For this, it turns out useful to rewrite labor market clearing as

$$\int_0^s \int_0^1 L_{s',x} dx ds' = H_s \text{ for all } s,$$

where H_s is the cumulative distribution function of labor endowments. This specification allows us to embed mass points as jumps in H_s .

Suppose now that $\psi_{k,x}/\psi_{s',x}$ is not quasi-concave in x for s' and consider the case where only skill s' is supplied:

$$H_s = \mathbb{I}_{s > s'} = \begin{cases} 0 & \text{if } s < s', \\ 1 & \text{if } s \geq s'. \end{cases}$$

Since $\psi_{k,x}/\psi_{s,x}$ is not quasi-concave, there exist $x_1 < x_2 < x_3$ such that

$$\frac{w_{s'}\psi_{k,x_1}}{\psi_{s',x_1}} > \frac{w_{s'}\psi_{k,x_2}}{\psi_{s',x_2}} < \frac{w_{s'}\psi_{k,x_3}}{\psi_{s',x_3}}.$$

Because only labor of type s' is being supplied, Euler's theorem in this case implies that w_s equals net output. Next note that net output is continuous in the allocation and in q , and hence Berge's maximum theorem applies and implies that equilibrium net output is continuous in q (and equilibrium allocations are upper hemicontinuous in q). Moreover, net output is also increasing in q (see proposition 8). Thus, the wage w_s is continuously increasing in q , and there exists a value for q such that

$$\frac{w_s \psi_{k,x_1}}{\psi_{s',x_1}}, \frac{w_s \psi_{k,x_3}}{\psi_{s',x_3}} > \frac{1}{q} > \frac{w_s \psi_{k,x_2}}{\psi_{s',x_2}},$$

which implies that X_k cannot be convex.

It remains to extend the result to labor endowments without mass points, which is a simple continuity argument. Net output is continuous in allocations, while the set of feasible allocations is continuous in the endowment cumulative density function H . Thus, by the maximum theorem, the set of equilibrium allocations is upper hemicontinuous in H . Since there is no equilibrium allocation generating a convex X_k under the endowment function $\mathbb{I}_{s>s'}$ considered above, we can construct a sequence of differentiable endowment functions with strictly positive derivative, $\{H^{(n)}\}_{n \in \mathbb{N}}$, that converges to $\mathbb{I}_{s>s'}$; for sufficiently large n , the set of automated tasks cannot be convex.

A2.3. Characterization of the Interior Automation Threshold

Here we characterize the productivity threshold q_m at which automation transitions from interior to low skill or high skill, respectively. For a characterization of the threshold q_0 at which automation starts, see the proof of proposition B1 in appendix B. We assume that our assumptions 1–3 and condition 1 are satisfied.

We distinguish three cases. First, suppose that condition 2 holds (as in part 3 of proposition 3). We know from condition 1 that $\psi_{k,x}/\psi_{0,x}$ is strictly increasing in x on a neighborhood of $x = 0$. Thus, we can define a threshold task \bar{x}_m as the smallest $x \in (0, 1)$ such that $\psi_{k,0}/\psi_{0,0} = \psi_{k,x}/\psi_{0,x}$. That is, the productivity ratio between capital and the least skilled workers is the same at task \bar{x}_m and task 0. Note that such an \bar{x}_m exists if and only if $\psi_{0,0}/\psi_{k,0} < \psi_{0,1}/\psi_{k,1}$, which is condition 2.1.

Now suppose that we restrict capital to tasks below \bar{x}_m and labor to tasks above \bar{x}_m . Then we choose the allocation that maximizes net output subject to these restrictions. Let w_s^m be the resulting wage function. Note that w_s^m is strictly increasing in q , allowing us to define q_m as the unique value of q that solves $1/q_m = w_0^m(q_m) \psi_{k,\bar{x}_m}/\psi_{0,\bar{x}_m}$, where we write $w_0^m(q)$ to emphasize the dependence of w_0^m on q . Intuitively, this condition equates the costs of producing task \bar{x}_m with capital and with the least skilled workers.

Finally, note that if $q = q_m$, the restriction of capital to tasks below \bar{x}_m and labor to tasks above \bar{x}_m is not binding, and in this case we have $X_k = [0, \bar{x}_m]$.

For the second case, suppose that both conditions 2.1 and 2.2 are violated strictly (as in part 5 of proposition 3). This case is completely symmetric to the first case: we define \bar{x}_m as the largest $x \in (0, 1)$ such that $\psi_{k,1}/\psi_{1,1} = \psi_{k,x}/\psi_{1,x}$ and q_m analogously to the first case.

In the third case, condition 2.1 is violated strictly, while condition 2.2 holds. Then, a threshold task \bar{x}_m as defined in either the first or the second case does not exist, and we simply set $q_m = q_\circ$.

A2.4. Proof of Proposition 3: Interior Automation

In appendix B, we establish proposition B1, which is a more general version of proposition 3. Compared with proposition B1, in proposition 3 we additionally impose assumption 3, which implies a convex set of automated tasks (see proposition 2). With a convex set of automated tasks, weakly interior automation implies interior automation, weakly low-skill automation implies low-skill automation, and weakly high-skill automation implies high-skill automation. Thus, parts 1, 2, 3, and 5 of proposition 3 follow immediately from proposition B1 and the fact that assumption 3 guarantees a convex set of automated tasks in proposition 3.

It remains to prove part 4 of proposition 3, which is the case in which the automation pattern is ambiguous under the assumptions of proposition B1. We prove part 4 by showing that (i) automation cannot be high skill under condition 2.2 and (ii) automation cannot be low skill if condition 2.1 is violated strictly. Together, this implies that automation is interior for all $q \in (q_0, q_\circ)$ if condition 2.1 is violated strictly while condition 2.2 holds, which is part 4 of proposition 3.

No high-skill automation under condition 2.2.—We first show that there cannot be high-skill automation if condition 2.2 holds. The proof is by contradiction.

Suppose $1 \in X_k$. Then, it must be cheaper to produce task $x = 1$ with capital than with labor, $w_1\psi_{k,1}/\psi_{s,1} \geq 1/q$ for all s . Moreover, this inequality must hold strictly for all but the most skilled workers $s = 1$, because assumption 2 implies that among all labor types, task $x = 1$ can be produced at the lowest cost using skill $s = 1$, $S_1^{\min} = \{1\}$.

At the same time, combining $w_1\psi_{k,1}/\psi_{1,1} \geq 1/q$ with condition 2.2 implies that it must be strictly cheaper to produce task $x = 0$ with capital than with the most skilled workers, $w_1\psi_{k,0}/\psi_{1,0} > 1/q$. By continuity of labor productivity $\psi_{s,x}$, this extends to some neighborhood of $s = 1$, that is, there exists $\epsilon > 0$ such that $w_s\psi_{k,0}/\psi_{s,0} > 1/q$ for all $s \in (1 - \epsilon, 1]$.

Hence, we have shown that for every $s \in (1 - \epsilon, 1)$, s is strictly more expensive than capital in both the most and the least complex task. Quasi-concavity of $\psi_{k,x}/\psi_{s,x}$ (assumption 3) requires that this extends to all tasks:

$$\frac{w_s\psi_{k,x}}{\psi_{s,x}} > \frac{1}{q} \text{ for all } s \in (1 - \epsilon, 1).$$

But this implies that skill levels $s \in (1 - \epsilon, 1)$ cannot be assigned to any task in equilibrium, which is clearly incompatible with an equilibrium allocation maximizing (finite) net output. Hence, we must have $1 \notin X_k$: the most complex task is not automated.

No low-skill automation without condition 2.1.—Now suppose that condition 2.1 is violated strictly, that is, $\psi_{0,0}/\psi_{k,0} > \psi_{0,1}/\psi_{k,1}$. Then, arguments entirely symmetric to those of the previous paragraph show that there cannot be low-skill automation, such that we must have $1 \notin X_k$.

A2.5. Proof of Proposition 4: Minimum Wages and Automation

The first part of the proof follows closely the proof of part 1 of proposition A1. In particular, condition (8) together with comparative advantage across labor types (assumption 2) implies that $\psi_{k,x}/\psi_{s,x}$ is decreasing in x for all $s > \underline{s}$ and, by continuity, also for $s = \underline{s}$ (where \underline{s} is such that all skills below \underline{s} are nonemployed because of the minimum wage). Thus, the minimum labor cost function ω_x is decreasing, and the set of automated tasks, which is equal to $\{x | \omega_x \geq 1/q\}$, is either empty or contains zero.

The second part is to show that if $q \in (q_0, q_m)$, then the set of automated tasks remains nonempty after the introduction of the minimum wage. For this, we compare the assignment problems without capital, with and without the minimum wage.

Without capital, our setting has been studied extensively in the literature (e.g., Costinot and Vogel 2010). The introduction of the minimum wage is equivalent to a shift in the lower bound of the skill space from zero to \underline{s} . Without capital, it leads to a decline in skill premia along the entire skill space, with the wage of the least skilled remaining worker type \underline{s} increasing and the wage of the most skilled worker type ($s = 1$) decreasing (Teulings 2000; Costinot and Vogel 2010). Let w_s^0 be the wage function without capital and without minimum wage (and ω_x^0 the associated minimum labor cost function), and let $w_s^{0,\min}$ be the wage function without capital but with minimum wage. Then,

$$\frac{1}{q} < \max_x \omega_x^0 \leq \max_x \frac{w_s^0 \psi_{k,x}}{\psi_{s,x}} \leq \frac{w_{\underline{s}}^{0,\min} \psi_{k,0}}{\psi_{\underline{s},0}},$$

where the first inequality uses that $q > q_0$, the second follows from the definition of ω_x^0 as the lower envelope of all workers' effective cost, and the last inequality is implied by $w_s^0 \leq w_s^{0,\min}$ and $\psi_{k,x}/\psi_{s,x}$ being decreasing in x . The inequalities imply that if X_k were empty, we had $0 \in X_k$ by lemma A1, a contradiction. Hence, X_k remains nonempty after introduction of the minimum wage.

A2.6. Proof of Proposition 5: Labor Supply and Automation

We start with a useful lemma on the wage effects of labor supply changes that holds for all settings where production is concave and linear homogeneous in labor and wages equal marginal products (the proof is presented in appendix B).

LEMMA A2. Consider any two labor endowments $l > 0$ and $l^{\text{new}} > 0$ with corresponding wage functions w and w^{new} . Then, if $w_s \leq w_s^{\text{new}}$ for all s , we must have $w = w^{\text{new}}$.

The important implication of lemma A2 is that a labor supply change alone can never cause all wages to increase or all wages to decrease. Instead, there will always be some wages that increase and some that decrease, except in the case where the wage function is completely unchanged.

We can now prove that the threshold where automation transitions from interior to low skill, q_m , is strictly decreasing in relative skill supply. First, recall from characterization of q_m in appendix section A2.3 that q_m is defined as the unique solution to

$$\frac{1}{q_m} = \frac{w_0^m(q_m)\psi_{k,\bar{x}_m}}{\psi_{0,\bar{x}_m}}, \quad (\text{A2})$$

where w^m is the wage function obtained under the restriction that capital can be allocated only to tasks below \bar{x}_m and labor only to tasks above \bar{x}_m . This equation has a unique solution because the left-hand side is strictly decreasing and the right-hand side strictly increasing in q_m .

Now consider an increase in relative skill supply, $\Delta \log l_s$ with $\Delta \log l_s$ strictly increasing in s . We know from prior work (e.g., Costinot and Vogel 2010) that in the pure assignment model without capital, such a change in labor supply would lower all skill premia. With the restriction that capital must be assigned below and labor above \bar{x}_m , the labor allocation is determined as in a pure labor assignment model. Hence, the result from prior work applies, and the wage change $\Delta \log w_s^m$ must be strictly decreasing in s . By lemma A2, the wage change cannot be negative for all skill levels, and we must have $\Delta \log w_0^m > 0$.²⁶ Thus, the right-hand side of equation (A2) increases such that q_m must decrease to solve equation (A2), so $\Delta q_m < 0$.

A3. Proofs for Section IV: Local Effects of Automation

We use our characterization of the wage and the assignment function in terms of the differential equation system (9)–(13) to conduct comparative statics with respect to capital productivity. Implicitly, this imposes assumptions 1–3 and $q \geq q_0$.

We consider a small change in capital productivity $d \log q$ (if $q = q_0$, we impose $d \log q > 0$ such that our equilibrium characterization continues to hold) and study its first-order effects on wages and assignment. From equations (9) and (11), we obtain the variational equations

$$(d \log w_s)' = \frac{\partial^2 \log \psi_{s,X}}{\partial s \partial x} dX_s, \quad (\text{A3})$$

$$(dX_s)' = \lambda \frac{l_s w_s^\lambda}{Y \psi_{s,X}^{\lambda-1}} d \log w_s - \frac{l_s w_s^\lambda}{Y \psi_{s,X}^{\lambda-1}} d \log Y - (\lambda - 1) \frac{l_s w_s^\lambda}{Y \psi_{s,X}^{\lambda-1}} \frac{\partial \log \psi_{s,X}}{\partial x} dX_s, \quad (\text{A4})$$

which hold for all $s \neq \tilde{s}$. The boundary conditions for the upper branch of these variations, that is, the branch on $(\tilde{s}, 1]$, are given by

$$\begin{aligned} d \log w_{\tilde{s}}^+ &= \left(\frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial x} - \frac{\partial \log \psi_{k,\bar{x}}}{\partial x} \right) d\bar{x} + \frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial s} d\tilde{s} - d \log q - (\log w_{\tilde{s}})^+ d\tilde{s} \\ &= \left(\frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial x} - \frac{\partial \log \psi_{k,\bar{x}}}{\partial x} \right) d\bar{x} - d \log q, \end{aligned} \quad (\text{A5})$$

$$\begin{aligned} dX_{\tilde{s}}^+ &= d\bar{x} - (X_{\tilde{s}})^+ d\tilde{s} \\ &= d\bar{x} - \frac{l_{\tilde{s}} w_{\tilde{s}}^\lambda}{Y \psi_{\tilde{s},\bar{x}}^{\lambda-1}} d\tilde{s}, \end{aligned} \quad (\text{A6})$$

²⁶ Note that the proof of lemma A2 uses only linear homogeneity and concavity of net output in labor, so it also applies to the situation where capital is restricted to tasks below and labor to tasks above \bar{x}_m .

where the superscript plus sign denotes the right-hand-side limit of the respective function. The boundary conditions for the lower branch (which exists only if automation is interior) are

$$\begin{aligned} d \log w_{\tilde{s}}^- &= \left(\frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} - \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x} + \frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{s}} d\tilde{s} - d \log q - (\log w_{\tilde{s}})^{-} d\tilde{s} \\ &= \left(\frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} - \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x} - d \log q, \end{aligned} \quad (\text{A7})$$

$$dX_{\tilde{s}}^- = d\tilde{x} - (X_{\tilde{s}})^{-} d\tilde{s}$$

$$= d\tilde{x} - \frac{l_k w_k^{\lambda}}{Y \psi_{\tilde{s}, \tilde{x}}^{\lambda-1}} d\tilde{s}, \quad (\text{A8})$$

with the superscript minus sign denoting left-hand-side limits.

From the upper branch of the system, we can obtain the change in assignment of the most skilled workers, $dX_l(d\tilde{x}, d\tilde{s})$, as a function of $d\tilde{x}$ and $d\tilde{s}$. Analogously, if automation is interior, the lower branch yields $dX_0(d\tilde{x}, d\tilde{s})$, the change in assignment of the least skilled workers as a function of $d\tilde{x}$ and $d\tilde{s}$. Both of these changes must be zero in equilibrium, which defines functions $d\tilde{x}(d\tilde{s})$ and $d\tilde{x}(d\tilde{s})$. The following lemma establishes some properties of $d\tilde{x}(d\tilde{s})$ and $d\tilde{x}(d\tilde{s})$.

LEMMA A3. Suppose that assumptions 1–3 hold, $q \geq q_0$, and $d \log q > 0$. Then, if $\tilde{x} < 1$, the function $d\tilde{x}(d\tilde{s})$ is strictly increasing and satisfies $d\tilde{x}(0) > 0$ and

$$\left(\frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} - \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x}(0) < d \log q + \frac{1}{\lambda} d \log Y.$$

Moreover, if $\tilde{x} > 0$, $d\tilde{x}(d\tilde{s})$ is strictly increasing and satisfies $d\tilde{x}(0) < 0$ and

$$\left(\frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} - \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x}(0) < d \log q + \frac{1}{\lambda} d \log Y.$$

The proof of lemma A3 is presented in appendix B.

A3.1. Proof of Proposition 6: Automation and Employment Polarization

Expansion of automation.—Suppose now that condition 1 holds and $q \in (q_0, q_m)$ such that automation is interior. In this case, condition (10) requires that

$$\underbrace{\left(\frac{\partial \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} + \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x}(d\tilde{s}) + \frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{s}} d\tilde{s}}_{=\gamma_{l, \tilde{x}}} = \underbrace{\left(\frac{\partial \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{x}} + \frac{\partial \log \psi_{k, \tilde{x}}}{\partial \tilde{x}} \right) d\tilde{x}(d\tilde{s})}_{=\gamma_{l, \tilde{x}}} + \frac{\partial \log \psi_{\tilde{s}, \tilde{x}}}{\partial \tilde{s}} d\tilde{s}, \quad (\text{A9})$$

where we have already inserted the functions $d\bar{x}(d\tilde{s})$ and $d\underline{x}(d\tilde{s})$ derived from equations (A3)–(A8). Rearranging and signing terms, we obtain

$$\underbrace{\gamma_{\tilde{s},\bar{x}} d\bar{x}(d\tilde{s})}_{\geq 0} - \underbrace{\gamma_{\tilde{s},\underline{x}} d\underline{x}(d\tilde{s})}_{\leq 0} = \underbrace{\left(\frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial s} - \frac{\partial \log \psi_{\tilde{s},\underline{x}}}{\partial s} \right) d\tilde{s}}_{<0}.$$

By lemma A3, the left-hand side of this equation is increasing while the right-hand side is strictly decreasing in $d\tilde{s}$. Thus, the equation determines a unique equilibrium change $d\tilde{s}^*$.

If $\gamma_{\tilde{s},\bar{x}} d\bar{x}(0) - \gamma_{\tilde{s},\underline{x}} d\underline{x}(0)$ is strictly positive, then $d\tilde{s}^*$ must be strictly negative, which implies that $d\underline{x}(d\tilde{s}^*) < 0$ (by lemma A3) and

$$d\bar{x}(d\tilde{s}^*) = \frac{\gamma_{\tilde{s},\bar{x}}}{\gamma_{\tilde{s},\underline{x}}} d\underline{x}(d\tilde{s}^*) + \frac{1}{\gamma_{\tilde{s},\bar{x}}} \left(\frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial s} - \frac{\partial \log \psi_{\tilde{s},\underline{x}}}{\partial s} \right) d\tilde{s}^* > 0.$$

Note here that $\gamma_{\tilde{s},\bar{x}} d\bar{x}(0) - \gamma_{\tilde{s},\underline{x}} d\underline{x}(0) > 0$ can hold only if $\gamma_{\tilde{s},\bar{x}} > 0$.

Analogously, if $\gamma_{\tilde{s},\bar{x}} d\bar{x}(0) - \gamma_{\tilde{s},\underline{x}} d\underline{x}(0)$ is strictly negative, then $d\tilde{s}^*$ must be strictly positive, which implies that $d\bar{x}(d\tilde{s}^*) > 0$ and $d\underline{x}(d\tilde{s}^*) < 0$.

Finally, if $\gamma_{\tilde{s},\bar{x}} d\bar{x}(0) - \gamma_{\tilde{s},\underline{x}} d\underline{x}(0)$ equals zero, then $d\tilde{s}^*$ must be zero as well such that $d\bar{x}(\tilde{s}^*) > 0$ and $d\underline{x}(d\tilde{s}^*) < 0$ follow immediately from lemma A3.

At this point, note also that

$$\max\{\gamma_{\tilde{s},\bar{x}} d\bar{x}(d\tilde{s}^*), \gamma_{\tilde{s},\underline{x}} d\underline{x}(d\tilde{s}^*)\} \leq \max\{\gamma_{\tilde{s},\bar{x}} d\bar{x}(0), \gamma_{\tilde{s},\underline{x}} d\underline{x}(0)\} < d \log q + \frac{1}{\lambda} d \log Y,$$

where the second inequality follows from lemma A3. This implies that the initial values $d \log \tilde{w}_s^+$ and $d \log \tilde{w}_s^-$ must both be negative at $d\tilde{s}^*$. This result will be useful in the next step of the proof.

Employment polarization.—We have just shown above that $d \log \tilde{w}_s^+$ is strictly negative in equilibrium. This implies that the initial value dX_s^+ in the dynamic system for dX_s and $d \log \tilde{w}_s$ must be strictly positive:

$$dX_s^+ = d\bar{x} - \frac{k_s w_s^\lambda}{Y \psi_{\tilde{s},\bar{x}}^{\lambda-1}} d\tilde{s} > 0.$$

If it were negative, we could never attain $dX_1 = 0$ by the reasoning in the proof of lemma A3.

Suppose now that at some skill $s_1 > \tilde{s}$, dX_s turns negative, that is, it crosses zero from above: $dX_{s_1} = 0$ and $(dX_s)' \leq 0$. To obtain $dX_1 = 0$, dX_s must at some point $s_2 > s_1$ attain zero again, this time from below: $dX_{s_2} = 0$ and $(dX_s)' \geq 0$. The differential equation for $(dX_s)'$, however, implies $(dX_s)' = \beta(s) d \log \tilde{w}_s$ for $s = s_1, s_2$. We must therefore have $d \log \tilde{w}_{s_1} \leq 0$. Since dX_s is negative between s_1 and s_2 , we will also have $d \log \tilde{w}_{s_2} < 0$ by the equation for $(d \log \tilde{w}_s)'$. This in turn implies $(dX_{s_2})' < 0$, a contradiction. As a result, dX_s cannot cross zero but stays positive until dX_1 .²⁷ Analogous reasoning yields $dX_s < 0$ for $0 < s < \tilde{s}$.

²⁷ We can exclude the case where dX_s has a critical zero (a point where dX_s is tangent to zero but does not cross it). This is because a critical zero would imply $(dX_s)' = 0$ and

Labor share.—Instead of proving the results for the labor share directly, we prove that the inverse of these results holds for the capital share. From equation (14), we obtain the response of the capital share to the increase in capital productivity as

$$\begin{aligned} d\alpha_k &= \alpha_k(\lambda - 1)d\log q + q^{\lambda-1} \frac{\partial \Gamma_k}{\partial \underline{x}} d\underline{x} + q^{\lambda-1} \frac{\partial \Gamma_k}{\partial \bar{x}} d\bar{x} \\ &= \alpha_k(\lambda - 1)d\log q - q^{\lambda-1} \psi_{k,\underline{x}}^{\lambda-1} d\underline{x} + q^{\lambda-1} \psi_{k,\bar{x}}^{\lambda-1} d\bar{x}. \end{aligned} \quad (\text{A10})$$

The last two terms are strictly positive by our employment polarization result such that the capital share increases if $\lambda \geq 1$.

If $\lambda < 1$, the total effect on the capital share depends on the relative strength of the capital deepening effect (first term) and the expansion of the set of automated tasks (second and third terms). If $q = q_0$, we have $\alpha_k = 0$ such that the capital deepening effect vanishes. Moreover, lemma A3 implies that

$$\max\{|d\bar{x}|, |d\underline{x}|\} \geq \min\{|d\bar{x}(0)|, |d\underline{x}(0)|\} > 0,$$

which means that the expansion of the set of automated tasks does not vanish. So, we must have $d\alpha_k(q_0) > 0$. Finally, note that $d\alpha_k$ (considering the perturbation $d\log q > 0$) is a right-hand derivative and, as such, it is continuous from the right, that is, $\lim_{q \searrow q_0} d\alpha_k = d\alpha_k(q_0) > 0$. This proves that $d\alpha_k > 0$ in some right neighborhood of q_0 .

A3.2. Proof of Proposition 7: Automation and Wage Polarization

By (A3), we have

$$d\log w_s - d\log w_{s'} = \int_s^{s'} \frac{\partial^2 \log \psi_{t,x}}{\partial s \partial x} dX_t dt.$$

By assumption 2 and our employment polarization result in proposition 6, this expression is strictly positive for all $s > s' \geq \tilde{s}$ and also for all $s < s' \leq \tilde{s}$.

A3.3. Proof of Proposition 8: Automation and Wage Levels

We have already proved part 1 of the proposition in the main text. Here we prove parts 2–4.

Part 2.—Condition (10) implies that

$$d\log w_{\tilde{s}} = \begin{cases} \frac{\partial \log \psi_{\tilde{s},x}}{\partial s} d\tilde{s} + \gamma_{\tilde{s},x} d\underline{x} - d\log q - (d\log w_{\tilde{s}})^{-} = \gamma_{\tilde{s},x} d\underline{x} - d\log q & \text{if } d\tilde{s} \geq 0, \\ \frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial s} d\tilde{s} + \gamma_{\tilde{s},\bar{x}} d\bar{x} - d\log q - (d\log w_{\tilde{s}})^{+} = \gamma_{\tilde{s},\bar{x}} d\bar{x} - d\log q & \text{if } d\tilde{s} \leq 0, \end{cases} \quad (\text{A11})$$

$dX_t = 0$. But then, the entire upper branch of dX_t would be identically zero, which is incompatible with the initial value dX_t^+ being strictly positive.

where $\gamma_{\tilde{s},x}$ and $\gamma_{\tilde{s},\bar{x}}$ are defined as in equation (A9). Now suppose at first that $q = q_0$. Then, $\bar{x} = \underline{x}$ and $\gamma_{\tilde{s},\underline{x}} = \gamma_{\tilde{s},\bar{x}} = 0$ because $\underline{x} = \bar{x}$ is a maximizer of the effective labor cost function ω_x . So, we obtain $d \log w_{\tilde{s}} = -d \log q < 0$.

Next, for $d \log q > 0$, $d \log w_{\tilde{s}}$ is a right-hand derivative and thus must be continuous from the right. So, $d \log w_{\tilde{s}} < 0$ in a right neighborhood of q_0 . Finally, this extends to skills in some neighborhood around \tilde{s} because the wage change $d \log w_s$ is continuous in s .

Part 3.—If $\psi'_{s',x}/\psi_{k,x}$ is constant, we must have $w_{s'}/\psi_{s',x} = 1/(q\psi_{k,x})$ for all x .²⁸ Differentiating this, we obtain $d \log w_{s'} = -d \log q < 0$. Since the change $d \log w_s$ is continuous in s , we obtain $d \log w_s < 0$ for all s in some neighborhood of s' .

Part 4.—Suppose at first that $q = q_m$ and consider $d \log q < 0$. It is easy to check that for $d \log q < 0$, the reasoning of lemma A3 can be adjusted to imply that $d\bar{x}(d\tilde{s})$ is still strictly increasing but now $d\bar{x}(0) < 0$ and

$$\left(\frac{\partial \log \psi_{\tilde{s},\bar{x}}}{\partial x} - \frac{\partial \log \psi_{k,\bar{x}}}{\partial x} \right) d\bar{x}(0) > d \log q + \frac{1}{\lambda} d \log Y.$$

Since at $q = q_m$ we have $\underline{x} = 0$ and $\tilde{s} = 0$, the analogous results for $d\underline{x}$ do not apply. Instead, we have $dX_{\tilde{s}}^- = 0$, and hence by equation (A8),

$$d\underline{x}(d\tilde{s}) = \frac{l_0 w_0^\lambda}{Y \psi_{\tilde{s},x}^{\lambda-1}} d\tilde{s}.$$

So $d\underline{x}(\tilde{s})$ is strictly increasing and $d\underline{x}(0) = 0$. Since we are considering $d \log q < 0$ starting from q_m , equation (10) holds, and so does its variational counterpart (A9). Then, by reasoning analogous to that in the first part of the proof of proposition 6, we can show that $d\bar{x} < 0$ and $d\underline{x} \geq 0$. Next, by the same reasoning as in the second part of the proof of proposition 6, we obtain that $dX_s < 0$ for all $s \in (0, 1)$. By the argument in the proof of proposition 7, this implies that $d \log w_l < d \log w_0$.

Finally, note that for $d \log q < 0$, the changes $d \log w_l$ and $d \log w_0$ are left-hand derivatives and, as such, they are continuous from the left. So, we have that $d \log w_l < d \log w_0$ in response to $d \log q < 0$ for all q in some left neighborhood of q_m . But for $q \in (q_0, q_m)$, wages are differentiable in q , and we obtain the reverse for $d \log q > 0$, that is, $d \log w_l > d \log w_0$ in response to $d \log q > 0$ for all q in some left neighborhood of q_m .

A3.4. Proof of Proposition 9: Productivity Effects

We first derive a second-order approximation of net output NY . Net output is given by $NY = \max_{\bar{K}} F(\bar{K}, l) - \bar{K}/q$, where F is maximal output subject to aggregate factor supplies \bar{K} and l (see proof of proposition 1). So by the envelope theorem, we obtain $dNY/dq = \bar{K}/q^2$, and hence $d \log NY / d \log q = \bar{K}/(qNY) = \bar{K}/(qY - \bar{K}) = \alpha_k/(1 - \alpha_k)$. The second-order term can then be written as

²⁸ Otherwise, either the set of automated tasks was empty (if the equation held with $<$ instead of $=$) or skill s' could not be assigned to any task (if the equation held with $>$ instead of $=$).

$$\frac{d^2 \log NY}{(d \log q)^2} = \frac{d}{d \log q} \frac{d \log NY}{d \log q} = \frac{d(\alpha_k / (1 - \alpha_k))}{d \log q} = \frac{\alpha_k}{(1 - \alpha_k)^2} \frac{d \log \alpha_k}{d \log q},$$

which by our previous result (A10) can be written as

$$\frac{\alpha_k}{(1 - \alpha_k)^2} \left(\lambda - 1 + \frac{\partial \log \Gamma_k}{\partial \underline{x}} \frac{d \underline{x}}{d \log q} + \frac{\partial \log \Gamma_k}{\partial \bar{x}} \frac{d \bar{x}}{d \log q} \right).$$

Combining these first- and second-order terms yields our second-order Taylor approximation:

$$\begin{aligned} \Delta \log NY \approx & \frac{\alpha_k}{1 - \alpha_k} \Delta \log q \\ & + \frac{\alpha_k}{(1 - \alpha_k)^2} \left[\lambda - 1 + \frac{\partial \log \Gamma_k}{\partial \bar{x}} \frac{d \bar{x}}{d \log q} + \frac{\partial \log \Gamma_k}{\partial \underline{x}} \frac{d \underline{x}}{d \log q} \right] (\Delta \log q)^2. \end{aligned}$$

Next, we translate this into an expression for TFP. Since $Y = NY + \bar{K}/q$, we have $\Delta \log Y \approx (1 - \alpha_k) \Delta \log NY + \alpha_k \Delta \log (\bar{K}/q)$. Using this in the definition of TFP, we obtain $\Delta \log TFP \approx (1 - \alpha_k) \Delta \log NY$. Plugging in the second-order approximation of $\Delta \log NY$ derived above, we obtain the expression for TFP from proposition 9.

Data Availability

Data and code replicating the quantitative analysis in section VI and in appendix B, including all figures and tables, can be found in Acemoglu and Loebbing (2025) in the Harvard Dataverse, <https://doi.org/10.7910/DVN/ATCOI7>.

References

Acemoglu, D. 1998. "Why Do New Technologies Complement Skills? Directed Technical Change and Wage Inequality." *Q.J.E.* 113 (4): 1055–89.

—. 1999. "Changes in Unemployment and Wage Inequality: An Alternative Theory and Some Evidence." *A.E.R.* 89 (5): 1259–78.

—. 2007. "Equilibrium Bias of Technology." *Econometrica* 75 (5): 1371–409.

—. 2024. "The Simple Macroeconomics of AI." *Econ. Policy* 39 (120): 1–46.

Acemoglu, D., and D. H. Autor. 2011. "Skills, Tasks and Technologies: Implications for Employment and Earnings." in *Handbook of Labor Economics*, vol. 4, edited by D. Card and O. Ashenfelter, 1043–171. Amsterdam: Elsevier.

Acemoglu, D., M. K. Jensen, and P. Restrepo. 2025. "Automation and Long-run Wages." Working paper.

Acemoglu, D., and S. Johnson. 2023. *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity* Public Affairs.

Acemoglu, D., and J. Loebbing. 2025. "Replication Data for: 'Automation and Polarization.'" Harvard Dataverse, <https://doi.org/10.7910/DVN/ATCOI7>.

Acemoglu, D., and P. Restrepo. 2018a. "Modeling Automation." *AEA Papers and Proc.* 108:48–53.

_____. 2018b. "The Race between Machine and Man: Implications of Technology for Growth, Factor Shares and Employment." *A.E.R.* 108 (6): 1488–542.

_____. 2020. "Robots and Jobs: Evidence from US Labor Markets." *J.P.E.* 128 (6): 2188–244.

_____. 2022. "Tasks, Automation, and the Rise in US Wage Inequality." *Economica* 90 (5): 1973–2016.

Acemoglu, D., and F. Zilibotti. 2001. "Productivity Differences." *Q.J.E.* 116 (2): 563–606.

Ales, L., C. Combemale, E. R. H. Fuchs, and K. S. Whitefoot. 2024. "How It's Made: A General Theory of the Labor Implications of Technical Change." Working paper.

Autor, D. H. 2014. "Polanyi's Paradox and the Shape of Employment Growth." Working Paper no. 20485, NBER, Cambridge, MA.

_____. 2015. "Why Are There Still So Many Jobs? The History and Future of Workplace Automation." *J. Econ. Perspectives* 29 (3): 3–30.

Autor, D. H., and D. Dorn. 2013. "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market." *A.E.R.* 103 (5): 1553–97.

Autor, D. H., F. Levy, and R. J. Murnane. 2003. "The Skill Content of Recent Technological Change: An Empirical Exploration." *Q.J.E.* 118 (4): 1279–333.

BEA (US Bureau of Economic Analysis). 2021a. "Fixed Assets Accounts Tables, Table 2.7: Investment in Private Fixed Assets, Equipment, Structures, and Intellectual Property Products by Type." Suitland, MD: BEA.

_____. 2021b. "National Income and Product Accounts, Table 6.8: Persons Engaged in Production by Industry." Suitland, MD: BEA.

Berg, A., E. F. Buffie, and L.-F. Zanna. 2018. "Should We Fear the Robot Revolution? (The Correct Answer Is Yes)." *J. Monetary Econ.* 97:117–48.

Clemens, M. A., E. G. Lewis, and H. M. Postel. 2018. "Immigration Restrictions as Active Labor Market Policy: Evidence from the Mexican Bracero Exclusion." *A.E.R.* 108 (6): 1468–87.

Costinot, A., and J. Vogel. 2010. "Matching and Inequality in the World Economy." *J.P.E.* 118 (4): 747–86.

Dauth, W., S. Findeisen, J. Suedekum, and N. Woessner. 2021. "The Adjustment of Labor Markets to Robots." *J. European Econ. Assoc.* 19 (6): 3104–53.

DiCecio, R. 2009. "Sticky Wages and Sectoral Labor Comovement." *J. Econ. Dynamics and Control* 33 (3): 538–53.

Eloundou, T., S. Manning, P. Mishkin, and D. Rock. 2024. "GPTs are GPTs: Labor Market Impact Potential of LLMs." *Science* 384 (6702): 1306–8.

Feng, A., and G. Graetz. 2020. "Training Requirements, Automation, and Job Polarisation." *Econ. J.* 130 (631): 2249–71.

Ford, M. 2015. *Rise of the Robots: Technology and the Threat of a Jobless Future*. New York: Basic.

Goos, M., A. Manning, and A. Salomons. 2009. "Job Polarization in Europe." *A.E.R.* 99 (2): 58–63.

Graetz, G., and G. Michaels. 2018. "Robots at Work." *Rev. Econ. and Statis.* 100 (5): 753–68.

Greenwood, J., Z. Hercowitz, and P. Krusell. 1997. "Long-Run Implications of Investment-Specific Technological Change." *A.E.R.* 87 (3): 342–62.

Grossman, G. M., E. Helpman, E. Oberfield, and T. Sampson. 2017. "Balanced Growth despite Uzawa." *A.E.R.* 107 (4): 1293–312.

Heckman, J. J., and G. Sedlacek. 1985. "Heterogeneity, Aggregation, and Market Wage Functions: An Empirical Model of Self-Selection in the Labor Market." *J.P.E.* 93 (6): 1077–125.

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

Hémous, D., M. Olsen, C. Zanella, and A. Dechezleprêtre. 2025. “Induced Automation Innovation: Evidence from Firm-Level Patent Data.” *J.P.E.* 133 (6): 1975–2028.

Hubmer, J. 2023. “The Race between Preferences and Technology.” *Econometrica* 91 (1): 227–61.

Humlum, A. 2021. “Robot Adoption and Labor Market Dynamics.” Working paper.

IFR (International Federation of Robotics). 2018. “World Robotics 2018.” Frankfurt: IFR.

Jackson, M. O., and Z. Kanik. 2020. “How Automation That Substitutes for Labor Affects Production Networks, Growth, and Income Inequality.” Working paper.

Jaimovich, N., I. Saporta-Eksten, H. Siu, and Y. Yedid-Levi. 2021. “The Macroeconomics of Automation: Data, Theory, and Policy Analysis.” *J. Monetary Econ.* 122:1–16.

Lindenlaub, I. 2017. “Sorting Multidimensional Types: Theory and Application.” *Rev. Econ. Studies* 84 (2): 718–89.

Loebbing, J. 2022. “An Elementary Theory of Directed Technical Change and Wage Inequality.” *Rev. Econ. Studies* 89 (1): 411–51.

Moll, B., L. Rachel, and P. Restrepo. 2022. “Uneven Growth: Automation’s Impact on Income and Wealth Inequality.” *Econometrica* 90 (6): 2645–83.

Ocampo, S. 2022. “A Task-Based Theory of Occupations with Multidimensional Heterogeneity.” Working paper.

Polanyi, M. 1966. *The Tacit Dimension*. New York: Doubleday.

Rosen, S. 1974. “Hedonic Prices and Implicit Markets: Product Differentiation in Pure Competition.” *J.P.E.* 82 (1): 34–55.

Sattiger, M. 1975. “Comparative Advantage and the Distributions of Earnings and Abilities.” *Econometrica* 43 (3): 455–68.

———. 1993. “Assignment Models of the Distribution of Earnings.” *J. Econ. Literature* 31 (2): 831–80.

Stokey, N. L. 2018. “Technology, Skill, and the Wage Structure.” *J. Human Capital* 12 (2): 343–84.

Teulings, C. N. 1995. “The Wage Distribution in a Model of the Assignment of Skills to Jobs.” *J.P.E.* 103 (2): 280–315.

———. 2000. “Aggregation Bias in Elasticities of Substitution and the Minimum Wage Paradox.” *Internat. Econ. Rev.* 41 (2): 359–98.

———. 2005. “Comparative Advantage, Relative Wages, and the Accumulation of Human Capital.” *J.P.E.* 113 (2): 425–61.

Teulings, C. N., and P. A. Gautier. 2004. “The Right Man for the Job.” *Rev. Econ. Studies* 71 (2): 553–80.

Tinbergen, J. 1956. “On the Theory of Income Distribution.” *Weltwirtschaftliches Archiv* 77 (2): 156–73.

Topkis, D. M. 1998. *Supermodularity and Complementarity*. Princeton, NJ: Princeton Univ. Press.

Wiener, N. 1950. *The Human Use of Human Beings: Cybernetics and Society*. Boston: Houghton Mifflin.

Zeira, J. 1998. “Workers, Machines, and Economic Growth.” *Q.J.E.* 113 (4): 1091–117.