

NBER WORKING PAPER SERIES

AUTOMATION: THEORY, EVIDENCE, AND OUTLOOK

Pascual Restrepo

Working Paper 31910

<http://www.nber.org/papers/w31910>

NATIONAL BUREAU OF ECONOMIC RESEARCH

1050 Massachusetts Avenue

Cambridge, MA 02138

November 2023, Revised December 2023

Prepared for the Annual Review of Economics. When citing this paper, please use the following: Restrepo, Pascual. 2023. "Automation: Theory, Evidence, and Outlook." Annual Review of Economics. Submitted. DOI 10.1146/annurev-economics-090523-113355. This review article builds on joint work with Daron Acemoglu. I thank Daron Acemoglu, Philippe Aghion, David Autor, Alvaro Contreras, Emin Dinlersoz, Eric Donald, Georg Graetz, David Hémous, Joachim Hubmer, Anders Humlum, Xavier Jaravel, Fredric Kong, Juan Lopez del Valle, Ezra Oberfield, Richard Rogerson, Anna Salomons, and Joonas Tuhkuri for helpful comments. I am particularly thankful to Daron Acemoglu and Richard Rogerson for commenting on various versions of this manuscript. I received support from the National Science Foundation (award 2049427). The views expressed herein are those of the author and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2023 by Pascual Restrepo. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

Automation: Theory, Evidence, and Outlook
Pascual Restrepo
NBER Working Paper No. 31910
November 2023, Revised December 2023
JEL No. E24,J20

ABSTRACT

This article reviews the literature on automation and its impact on labor markets, wages, factor shares, and productivity. I first introduce the task model and explain why this framework offers a compelling way to think about recent labor market trends and the effects of automation technologies. The task model clarifies that automation technologies operate by substituting capital for labor in a widening range of tasks. This substitution reduces costs, creating a positive productivity effect, but also reduces employment opportunities for workers displaced from automated tasks, creating a negative displacement effect. I survey the empirical literature and conclude that there is wide qualitative support for the implications of task models and the displacement effects of automation. I conclude by discussing shortcomings of the existing literature and avenues for future research.

Pascual Restrepo
Department of Economics
Boston University
270 Bay State Road
Boston, MA 02215
and NBER
pascual@bu.edu

Over the past 40 years, the US economy has witnessed a notable decline in the labor share, particularly in major sectors like manufacturing and retail trade, a significant shift away from routine jobs in both factories and offices, and a sizable increase in wage inequality, with wages for workers without college degrees stagnating. While some of these trends are most pronounced in the US, they are visible in other advanced economies.¹

One explanation for these trends emphasizes the role of automation technologies: advances that enable substituting capital for labor at a widening range of tasks or processes. Examples include the developments of robotics enabling the substitution of robots for workers in manufacturing, the development of computer-numerically controlled machines, eliminating the need for machine operators, and the development of software systems that automate clerical tasks, such as handling payroll, logistics, and sales.²

This article reviews the literature on automation and its effects on labor markets and the economy. My point of departure is the *task model of automation*, in Section 1. The task model adopts the perspective that producing goods and services requires completing tasks. To produce a car, one has to design it, procure parts, assemble them, weld them, and so on. Tasks are assigned to groups of workers with different skills. But increasingly, automation technologies allow firms to produce their tasks using software systems, dedicated machinery, or industrial robots instead of workers. This displaces workers from automated tasks—the displacement effect—but lowers their cost—the productivity effect.

The task model highlights the distinct implications of automation for workers, firms, industries, and the occupational and wage structure:

- Automation reduces the labor share of adopting firms and industries.
- Automation shifts the occupational structure of firms, industries, and the economy by reducing labor demand in exposed occupations (those in which workers perform tasks that become automated).
- Automation reduces the relative demand for groups of workers who performed automated tasks via its displacement effect.
- Automation can reduce real wages and employment for displaced workers if they

¹For a summary of empirical trends in the US occupational and wage structure see [Goldin and Katz \(2008\)](#), [Acemoglu and Autor \(2011\)](#), and [Autor \(2019\)](#). For the labor share, see [Karabarbounis and Neiman \(2013\)](#), [Dao et al. \(2019\)](#), [Grossman and Oberfield \(2021\)](#), and [Hubmer and Restrepo \(2021\)](#).

²Historically, we have also had the automation (or mechanization) of tasks performed by workers. For example, the development of threshing machines in agriculture substituted for farm labor, while spinning and weaving machines substituted for skilled artisans during the Industrial Revolution.

cannot reallocate to non-automated tasks and its productivity gains are modest.

The task model also explains that these implications are unique to automation, distinguishing it from other forms of technological progress that do not displace workers from their tasks. These results single out automation as a potential driver of the labor share decline and the shifts in occupational and wage structures seen since the 1980s.

Section 2 turns to the empirical literature. This literature traces the development and diffusion of specific automation technologies (most notably, the introduction of industrial robots in manufacturing) and explores their impact on labor markets, the demand for skills, employment, factor shares, and productivity. This literature finds qualitative support for the implications of task models and the displacement effects of automation.

This section reviews the growing literature using firm data to explore how the adoption of new automation technology affects firms. Most work using firm-level data reports a reduction in adopting firms' labor shares and a shift in their workforce composition. Papers in this literature also report increased sales and employment among adopting firms relative to competitors. The theory section clarifies that firm-level employment expansions are to be expected but are not informative of the aggregate impacts of automation.

Section 3 concludes by discussing limitations of this literature and areas for future work. I emphasize the need to improve existing measures of automation and go beyond reduced-form evidence.

1 THE TASK MODEL

This section introduces the task model. My treatment follows [Acemoglu and Restrepo \(2022\)](#) and extends this by introducing firms to connect with the empirical literature.³

The Framework

A final good y is produced from differentiated products y_n with $n \in \mathcal{N} = \{1, \dots, N\}$. Products are combined using a constant-returns to scale technology $y = f(\{y_n\}_{n \in \mathcal{N}})$. Depending on the application, products represent industries or firms.

³The idea of modeling the substitution of capital for labor at the task level is from [Autor et al. \(2003\)](#). The model in [Acemoglu and Restrepo \(2022\)](#) builds on [Zeira \(1998\)](#), [Grossman and Rossi-Hansberg \(2008\)](#), [Acemoglu and Autor \(2011\)](#), [Acemoglu and Restrepo \(2018\)](#), and [Aghion et al. \(2018\)](#). For complementary approaches see [Jackson and Kanik \(2020\)](#), [Martinez \(2021\)](#), [Ocampo \(2022\)](#), and [Hémous and Olsen \(2022\)](#).

Each product requires completing a mass 1 of tasks x from disjoint sets \mathcal{T}_n .⁴ Task quantities y_x are aggregated with a constant elasticity of substitution $\lambda \geq 0$

$$y_n = \left(\int_{x \in \mathcal{T}_n} y_x^{\frac{\lambda-1}{\lambda}} \cdot dx \right)^{\frac{\lambda}{\lambda-1}}.$$

Tasks can be produced using labor or task-specific capital—software and equipment designed for this process. Workers belong to skill groups $g \in \{1, \dots, G\}$. Depending on the application, groups can be interpreted as workers with the same skills, observable attributes (i.e., education and age), or in a region. The total quantity of task x produced is

$$y_x = \sum_g \psi_{gx} \cdot \ell_{gx} + \psi_{kx} \cdot k_x,$$

where ℓ_{gx} is the quantity of labor allocated to task x , k_x is the quantity of task-specific capital in use, and $\psi_{gx} \geq 0$ and $\psi_{kx} \geq 0$ denote their productivity in task x . The ψ s vary by tasks and groups and encode their comparative advantage.

Task-specific capital and workers are perfect substitutes in the production of task x . This feature of the model captures in a stark way the fact that, in all instances of automation listed in the introduction, we have a machine or software that can, for all practical purposes, perfectly substitute for workers at narrowly defined tasks. A welding robot is a perfect substitute for humans in the task of welding car parts. A software system is a perfect substitute for humans in the task of receiving and dispatching sales orders.⁵

Capital k_x is produced from the final good at a constant unit cost $1/q_x$, where q_x is the efficiency with which the investment sector produces this capital. The remaining output is used for consumption c . Capital is produced and used each instant so that the economy's resource constraint is⁶

$$c + \sum_n \int_{x \in \mathcal{T}_n} (k_x/q_x) \cdot dx \leq y.$$

⁴Assuming disjoint sets eases notation and is without loss of generality since tasks can be relabeled (assembling product n can be labeled as a different task from assembling product n').

⁵This can be relaxed in two ways: one could assume automation is partial and capital must be combined with other workers to produce task x (e.g., spinning and weaving machines tended by women and children replacing skilled artisans, or software systems maintained by engineers substituting for workers in clerical positions). One could also allow for differentiation in tasks produced by capital and labor. This is relevant for customer interaction tasks, where some might prefer dealing with humans.

⁶This can also be viewed as the steady state of an economy that accumulates task-specific capital over time and faces a constant interest rate in the long run, with the constant user cost of capital folded in q_x .

Group g 's labor supply is $\ell_g = m_g \cdot w_g^\varepsilon$ with $\varepsilon \geq 0$ and w_g group's g wage. This labor supply can result from households' optimization over consumption and leisure or labor-market frictions as in [Kim and Vogel \(2021\)](#). Labor-market clearing for g requires

$$\sum_n \int_{x \in \mathcal{T}_n} \ell_{gx} \cdot dx = m_g \cdot w_g^\varepsilon.$$

Firms and industries pay the same wages w_g . There is no monopsony power or rents.⁷

A *competitive equilibrium* is given by a wage vector $w = \{w_g\}$ and a product price vector $p = \{p_n\}$ such that markets clear and the task allocation minimizes costs. [Acemoglu and Restrepo \(2022\)](#) provide conditions for the existence and uniqueness of equilibrium and conditions under which each task is assigned to a unique factor (except for a zero-measure indifference set). I assume these conditions hold so that the equilibrium is a partition of tasks into those assigned to capital and different workers, as in [Figure 1](#).




FIGURE 1: TASK ASSIGNMENT, TASK SHARES, AND AUTOMATION. The figure represents the tasks needed to complete different products $n = 1, \dots, N$ and how these tasks are assigned to workers of different skills (g) or capital (k). The gray areas represent newly automated tasks and the displacement effects.

Representing the equilibrium in terms of task shares

Equilibrium outcomes depend on the schedules of workers and capital productivities across tasks $\langle \{\psi_{gx}\}_g, \psi_{kx}, q_x \rangle$, which determine workers' comparative advantage and guide the assignment of tasks to factors. Previous theoretical work proposed specific parameterizations of these schedules to derive properties of the equilibrium and comparative statics. Recent

⁷See [Acemoglu and Restrepo \(2023\)](#) for studies exploring the role of rents in task models.

work by [Acemoglu and Restrepo \(2019\)](#) and [Acemoglu and Restrepo \(2022\)](#) developed a general approach to characterize the equilibrium in terms of *task shares*.

The task shares of capital and group g in product n are defined as

$$\Gamma_{nk}(w) = \int_{\mathcal{T}_{nk}(w)} (\psi_{kx} \cdot q_x)^{\lambda-1} \cdot dx, \quad \Gamma_{ng}(w) = \int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx,$$

where $\mathcal{T}_{nk}(w)$ and $\mathcal{T}_{ng}(w)$ denote the set of product- n tasks assigned to capital and labor when wages are $w = \{w_g\}$. Task shares are functions of technology and wages. They capture the importance of tasks assigned to capital and workers at each wage level and summarize all information on the assignment of tasks relevant to equilibrium outcomes.

All equilibrium objects can be computed in terms of task shares. For example, equilibrium wages $w = \{w_g\}$, good prices $p = \{p_n\}$, and output y solve the system of equations:

- the price of the final good is 1,

$$(1) \quad c^f(p) = 1,$$

where $c^f(p)$ is the unit cost function associated with f (its dual);

- p_n equals the marginal cost of producing y_n

$$(2) \quad p_n = \left(\Gamma_{nk}(w) + \sum_g \Gamma_{ng}(w) \cdot w_g^{1-\lambda} \right)^{\frac{1}{1-\lambda}};$$

- the labor market for g workers clears

$$(3) \quad w_g = \left(\frac{y}{m_g} \right)^{\frac{1}{\lambda+\varepsilon}} \cdot \left(\sum_n s_y^n(p) \cdot p_n^{\lambda-1} \cdot \Gamma_{ng}(w) \right)^{\frac{1}{\lambda+\varepsilon}},$$

with $s_y^n(p) = \frac{\partial \ln c^f(p)}{\partial \ln p_n}$ the share of product n in expenditure (Shephard's lemma).

Moreover, the equilibrium output of product n can be written as a constant elasticity of substitution (CES) production function, with task shares appearing as endogenous weights

$$(4) \quad y_n = \left(\Gamma_{nk}(w)^{\frac{1}{\lambda}} \cdot k_n^{\frac{\lambda-1}{\lambda}} + \sum_g \Gamma_{ng}(w)^{\frac{1}{\lambda}} \cdot \ell_{ng}^{\frac{\lambda-1}{\lambda}} \right)^{\frac{\lambda}{\lambda-1}}.$$

Relative to the usual CES, the novel aspect here is that the weight parameters that govern

the distribution of income are endogenous and respond to technology and wages. This allows task models to capture rich substitution patterns, including substitution between tasks (governed by λ) and within tasks (governed by the derivatives of task shares). This also allows for the possibility that technology, and in particular automation technologies, impact production by shifting these weights.

Modeling Automation

The examples show that automation technologies operate at the *extensive margin*: they substitute capital for labor in tasks that can now be automated but used to be performed by labor, displacing workers from these tasks. The simplest way to capture this process is by having $q_x = 0$ for some tasks assigned to workers initially. One can model the arrival of new automation technology as an exogenous increase in q_x from zero to $q'_x > 0$ for a specific set of tasks \mathcal{A} that can now be automated. To simplify the exposition, I assume that q'_x is large so that all tasks in \mathcal{A} are automated.⁸ This process is illustrated in Figure 1.

As shown in [Acemoglu and Restrepo \(2022\)](#), one can summarize the impact of this technology by two sufficient statistics: the *direct task displacement* it generates across groups, and the *cost-saving gains* from automating these tasks.

The direct task displacement on group g in product n is

$$d \ln \Gamma_{ng}^d = \frac{\int_{\mathcal{A}_{ng}} \psi_{gx}^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx},$$

where $\mathcal{A}_{ng} = \mathcal{A} \cap \mathcal{T}_{ng}$ denotes the set of tasks produced by workers of skill g in product n that become automated, and $d \ln \Gamma_{ng}^d$ gives the percent reduction in g 's task share from automating these tasks. This is computed at initial wages and does not account for the endogenous reassignment of tasks in response to wage changes.

The cost-saving gains from automating task x in \mathcal{A}_{ng} (at initial wages) is

$$\pi_x = \frac{1}{1-\lambda} \cdot \left(1 - \left[\frac{w_g \cdot \psi_{kx} \cdot q'_x}{\psi_{gx}} \right]^{\lambda-1} \right).$$

This is positive whenever the cost of producing the task with labor, w_g/ψ_{gx} , exceeds the cost of producing it with the new capital $1/(\psi_{kx} \cdot q'_x)$, which holds for all tasks in \mathcal{A}_{ng} by

⁸This can be extended to allow for a costly adoption margin (as in [Hubmer and Restrepo, 2021](#)) or by endogenizing advances as resulting from R&D investments (as in [Acemoglu and Restrepo, 2018](#)).

assumption.⁹ I let $\pi_{ng} > 0$ denote the cost-saving gains per task from automating tasks in \mathcal{A}_{ng} , defined as an (employment-weighted) average of π_x over these tasks.

The direct task displacement $d \ln \Gamma_{ng}^d$ describes how automation technology shifts task shares on impact, which depends on its capabilities to substitute for workers in different tasks.¹⁰ The cost-saving gains π_{ng} tells us how the productivity gains from automation vary depending on the technology and context. Automation technologies can bring large cost-saving gains if they are highly productive and the skills replaced are scarce, or small cost-saving gains otherwise (what [Acemoglu and Restrepo, 2019](#), call “so-so” technologies).

Equation (4) helps illustrate this notion of automation and how it reshapes the production process. Automation corresponds to an exogenous shift in task shares, expanding the CES weight of capital and reducing the CES weights of workers whose tasks were automated by $d \ln \Gamma_{ng}^d$. Automation can thus be conceived as a shift in the production process, placing more weight on capital and a lower weight on displaced workers.¹¹

Implications of Automation for Labor Shares, Occupations, and Workers

The next formulas summarize the impact of the exogenous arrival of a new automation technology. The formulas provide first-order approximations valid when the measure of automated tasks \mathcal{A} is small and are written in terms of the sufficient statistics $\{d \ln \Gamma_{ng}, \pi_{ng}\}$.

I provide two types of formulas. The first considers *direct effects*: the impact of automation on firm-level or industry-level outcomes, holding wages, aggregate output y , and aggregate product-price indices constant. These formulas give the differential effect of automation technologies on adopting firms’ outcomes relative to (otherwise comparable) non-adopting firms or industries and connect the theory to empirical work reporting cross-sectional estimates using firm and industry data.¹² The second considers *general equilibrium*

⁹ π_x can be approximated as $\pi_x \approx \ln(w_g/\psi_{gx}) - \ln(1/(\psi_{kx} \cdot q'_x))$ —the percent reduction in the cost of task x when automated. This approximation is valid when $\lambda \rightarrow 1$ or $w_g/\psi_{gx} \rightarrow 1/(\psi_{kx} \cdot q'_x)$ from above.

¹⁰Different waves of automation cause distinct shifts in task shares depending on their capabilities. Task models can thus explain why automation in the past was “de-skilling” (skilled artisans lost tasks to spinning and weaving machines) while it became “polarizing” in recent years (workers in middle-pay jobs lost tasks to robotics and software systems).

¹¹This notion of automation is missing from previous approaches, including work by [Krusell et al. \(2000\)](#) on capital-skill complementarity. These approaches start from production functions of the form $y = g(A_k \cdot k, \{A_g \cdot \ell_g\})$, where g is a CES (or nested CES) with exogenous weights and represent technology by an *increase* in some or all A ’s (i.e., factor-augmenting improvements) or as an increase in k due to greater investment efficiency. These models capture some aspects of technology but miss the shifting-weight role of automation. One exception is the work by [Zuleta \(2008\)](#) and [Peretto and Seater \(2013\)](#), who emphasize the possibility that technology shifts capital elasticities in a Cobb-Douglas production function.

¹²Section E in the Supplementary Materials provides the details for this connection. The idea is that firm outcomes can be written as a function of output, a vector of aggregate product-price indices, wages, and

effects on output and on group wages and employment levels. These formulas connect to empirical work that estimates the impact of automation on groups of workers (defined by skill, demographics, or regions). I denote direct effects by δa and GE effects by da .

Labor shares. I refer to $p_n \cdot y_n$ interchangeably as “sales” and “value added” since there are no other intermediates and to $s_n^\ell = \frac{\sum_g w_g \cdot \ell_{ng}}{p_n \cdot y_n}$ as the labor share (in both sales and value added). The direct effect of automation on firm or industry n labor share is

$$(5) \quad \delta \ln s_n^\ell = - \underbrace{\sum_g \omega_n^g \cdot d \ln \Gamma_{ng}^d}_{\text{displacement effects}} + \underbrace{(1 - \lambda) \cdot s_n^\ell \cdot \sum_g \omega_n^g \cdot d \ln \Gamma_{ng}^d \cdot \pi_{ng}}_{\text{task-price effects}}, \quad \left(\text{with } \omega_n^g = \frac{w_g \cdot \ell_{ng}}{\sum_j w_j \cdot \ell_{nj}} \right).$$

The displacement effect captures the negative impact of automation on the labor share due to the extensive margin reallocation of tasks from labor to capital. The task-price effect results from the reduction in the price of automated tasks. When $\lambda < 1$ so that tasks are complements, the cheaper automated tasks account for a smaller share of costs, raising the labor share. When $\lambda > 1$ so that tasks are substitutes, the cheaper automated tasks account for a higher share of costs, further reducing the labor share.

The displacement effect dominates for all values of λ , and automation reduces labor shares for adopting firms and industries relative to others. Section E shows that, in general, this also translates into an increase in sales per worker at adopting firms and industries.¹³

One important point is that the displacement effect from automation reduces adopters’ labor share independently of whether the elasticity of substitution between capital and labor is above or below one. In the task model, this elasticity (for product n) is

$$\sigma_n = \lambda + \underbrace{\frac{1}{1 - s_n^\ell} \cdot \sum_g \omega_n^g \cdot \left(- \sum_j \frac{\partial \ln \Gamma_{ng}(w)}{\partial \ln w_j} \right)}_{\text{substitution within marginal tasks } (\geq 0)}.$$

This exceeds λ , since firms can also substitute labor for capital at marginal tasks (the second term above). But σ_n might well be below one, depending on parameters. The

product n task shares. Well-identified empirical studies compare units with equal exposures to changes in output, product-price indices, and factor prices, all necessary for parallel trends. These studies identify the direct effects of automation working through changes in product n technology holding aggregates constant.

¹³Not all technological advances benefiting a firm increase sales per worker, even though this is commonly used as a measure of productivity. An increase in sales per worker relative to a competitor facing the same wages means that the firm is becoming more capital or skill intensive, not that it has become more productive.

reason why the effects of automation on the labor share are disconnected from σ_n is that elasticities of substitution summarize the impact of changes in input prices on firms' factor shares but are silent about the effects of shifts in task shares at the extensive margin. The task model thus reconciles the evidence for a negative labor share impact of automation reported below with studies that estimate $\sigma_n < 1$ (i.e., [Oberfield and Raval, 2020](#)).

A second important point is that the formula in (5) provides direct effects. The general equilibrium effects of automation on the aggregate labor share are derived in the appendix and are harder to sign. Besides its direct negative effect on the labor share of adopting firms and industries, automation affects the aggregate labor share by raising wage levels (as in [Grossman and Oberfield, 2021](#)) and reallocating economic activity across firms (as in [Oberfield and Raval, 2020](#)) or industries (as in [Acemoglu and Restrepo, 2022](#)).¹⁴

Occupational structure. While the task model is described in terms of tasks, employment and wage data are often collected and reported in terms of occupations. A simple way of thinking about occupations is as a partition of tasks, with each occupation $o \in \mathcal{O}$ defined as a bundle of tasks.

Denote by $\Gamma_{ong}(w)$ the task share of group g workers in occupation o tasks in product n , and by $d \ln \Gamma_{ong}^d$ the percent reduction in $\Gamma_{ong}(w)$ from automating tasks in \mathcal{A} . Differences in $d \ln \Gamma_{ong}^d$ capture the presence of tasks with varying potential for automation across occupations. For example, middle-pay occupations such as clerical or production jobs involve routine tasks that are easier to codify and automate ([Autor et al., 2003](#)). If recent automation technologies substitute for workers in routine tasks, the reduction in task shares $d \ln \Gamma_{ong}^d$ will concentrate on occupations involving routine tasks.

Let ω_n^o denote the share of wage payments made to workers in occupation o as a share of wages paid in n . The direct impact of automation on firm or industry n occupational wage shares is

$$(6) \quad \delta \ln \frac{\omega_n^o}{\omega_n^{o'}} = \underbrace{\sum_g \omega_{o'n}^g \cdot d \ln \Gamma_{o'ng}^d}_{\text{exposure to automation, } o'} - \underbrace{\sum_g \omega_{on}^g \cdot d \ln \Gamma_{ong}^d}_{\text{exposure to automation, } o}, \quad \left(\text{with } \omega_{on}^g = \frac{w_g \cdot \ell_{ong}}{\sum_j w_j \cdot \ell_{onj}} \right).$$

¹⁴These equilibrium effects are likely small for automation, with available estimates pointing to a weak increase in wage levels and a modest reallocation of economic activity across sectors in response to this shock. For example, in their quantitative exercise, [Acemoglu and Restrepo \(2022\)](#) estimate that automation increased wage levels by 6% and had a small impact on the sectoral composition of the economy since 1980. For firms, [Oberfield and Raval \(2020\)](#) find small differences in factor intensities across plants in an industry, implying a modest contribution from reallocation across firms to aggregate labor share changes.

This means that automation reduces firm and industry shares of wage payments (and employment) in occupations that are highly exposed to the technology, as measured by the reduction in task shares for workers in this occupation, $\sum_g \omega_{on}^g \cdot d \ln \Gamma_{ong}^d$.

GE effects on group-level outcomes, output, and product prices. Consider a shock directly changing group g wages (the right-hand-side of 3) by $\frac{1}{\lambda + \varepsilon} \cdot z_g$. In equilibrium, this shock leads to a reassignment of tasks, creating a fixed-point problem:

$$d \ln w_g = \underbrace{\frac{1}{\lambda + \varepsilon} \cdot z_g}_{\text{shocks}} + \underbrace{\frac{1}{\lambda + \varepsilon} \cdot \mathcal{J}_\Gamma \cdot d \ln w}_{\text{task reassignment}}$$

The reassignment of tasks is governed by the Jacobian \mathcal{J}_Γ , a $G \times G$ matrix where entry (g, j) is the elasticity of $\sum_n s_y^n(p) \cdot p_n^{\lambda-1} \cdot \Gamma_{ng}(w)$ with respect to w_j (holding p constant).

The solution to this fixed-point problem is

$$d \ln w_g = \Theta_g \cdot \text{stack}_j(z_j),$$

where $\text{stack}_j(z_j)$ is the column vector (z_1, z_2, \dots, z_G) . The change in group g wages depends on the vector of shocks experienced by all other groups and mediated by the *propagation matrix* $\Theta = \frac{1}{\lambda + \varepsilon} \left(\mathbb{I} - \frac{1}{\lambda + \varepsilon} \cdot \mathcal{J}_\Gamma \right)^{-1}$. This matrix has positive entries $\theta_{gj} \geq 0$, giving the extent to which a shock reducing the demand for j lowers g wages via the reassignment of tasks. This can be because j competes against g for tasks, or because j competes against other groups that compete with g , and so on. Θ is a Leontief inverse accumulating these effects.

The propagation matrix tells us how easily workers can reallocate and how this affects the incidence of labor demand shocks. If workers cannot easily reallocate, the propagation matrix will have large diagonal entries relative to off-diagonal ones and exposed groups will bear most of the incidence of shocks affecting them. If workers can easily reallocate, the propagation matrix will have small diagonal entries relative to off-diagonal ones and exposed groups will share the incidence of shocks with competing workers.

Using this matrix, one can compute the general equilibrium effects of automation on group outcomes, output, and product prices by solving a linear system of equations.

First, we have an equation for wage changes, obtained by differentiating (3):

$$(7) \quad d \ln w_g = \Theta_g \cdot \text{stack}_j \left(\underbrace{d \ln y}_{\text{productivity effect}} + \underbrace{\sum_n (\omega_j^n - s_y^n) \cdot d \ln \zeta_n}_{\text{change in product mix}} - \underbrace{\sum_n \omega_j^n \cdot d \ln \Gamma_{nj}^d}_{\text{displacement effects}} \right).$$

Here, $\omega_g^n = \ell_{ng}/\ell_g$ is the share of group g workers employed in n . The equilibrium effects on employment have a similar expression, computed from $d \ln \ell_g = \varepsilon \cdot d \ln w_g$.

Second, output changes are pinned down by the dual version of Solow's residual:

$$(8) \quad \sum_g s_y^g \cdot d \ln w_g = \sum_n s_y^n \cdot d \ln t f p_n \quad \left(\text{with } s_y^g = \frac{w_g \cdot \ell_g}{y} \right),$$

where $d \ln t f p_n$ denotes the increase in total factor productivity in product n :

$$d \ln t f p_n = s_n^\ell \cdot \sum_g \omega_n^g \cdot d \ln \Gamma_{ng}^d \cdot \pi_{ng} \geq 0.$$

This is positive in an efficient economy and depends on the cost-saving gains from automation $\pi_{ng} > 0$. Equation (8) shows that mean wage changes across groups add up to the aggregate TFP gains from automation. This holds in any competitive model with constant returns to scale and a fully elastic supply of capital.¹⁵

Third, we have an expression for the product shifters $d \ln \zeta_n$:

$$(9) \quad d \ln \zeta_n = (\lambda - 1) \cdot d \ln p_n + J_F \cdot d \ln p, \quad \text{with } d \ln p_n = s_n^\ell \cdot \sum_g \omega_n^g \cdot d \ln w_g - d \ln t f p_n.$$

Here, the Jacobian J_F is an $N \times N$ matrix whose entry (n, m) is the elasticity of s_y^n with respect to a change in p_m . For example, if F is a CES aggregator across products with an elasticity of substitution η , $s_y^n \propto p_n^{1-\eta}$ and J_F is a matrix with $1 - \eta$ in its diagonal.

Equations (7), (8), and (9) form a system of equations for the change in wages, prices, and output. From these equations, one can compute the effects of automation on group-level outcomes as a function of its sufficient statistics $\{d \ln \Gamma_{ng}^d, \pi_{ng}\}$.

Equation (7) bears particular relevance as it summarizes the channels through which automation affects wages and employment for group g .

- First, we have a positive productivity effect captured by the output expansion $d \ln y$.

¹⁵The idea that automation (and technology) increases mean wages in an efficient economy with an elastic capital supply goes back to [Simon \(1965\)](#) and was recently studied in [Caselli and Manning \(2019\)](#) and [Moll et al. \(2022\)](#).

This is positive and determined by the productivity gains from automation, $\{\pi_{ng}\}$.

- Second, we have changes in product mix, captured by the shifters $d \ln \zeta_n$. These capture the reallocation of economic activity towards firms and industries with different skill intensities. For example, automation could reallocate activity from manufacturing to high-skill service industries, increasing relative demand for skilled labor.¹⁶
- Third, and key to understanding the impact of automation, we have the task displacement effects on g workers, $\sum_n \omega_g^n \cdot d \ln \Gamma_{ng}^d$. Displacement effects reduce relative wages for exposed groups by leaving them with fewer employment opportunities.
- Fourth, general equilibrium effects depend on workers’ ability to reallocate and the substitution patterns summarized by the propagation matrix Θ .

To illustrate the implications of the displacement effects, consider a case with no product shifters and where the rows of the propagation matrix Θ add up to a common value so that all workers benefit equally from the productivity effect and changes in product mix.

Equation (7) shows that an automation shock displacing g workers reduces this group’s wage and employment relative to non-exposed groups. This is because the propagation matrix satisfies $\theta_{gg} > \theta_{jg}$ for $j \neq g$. This makes intuitive sense: workers displaced from some of their tasks are left with fewer employment opportunities than other groups. They can only reallocate by taking a (relative) wage cut. Exposed groups’ relative wage and employment decline is more pronounced when they cannot easily reallocate so that the propagation matrix is more diagonal and θ_{gg} is high relative to θ_{jg} .

Equation (7) also shows that automation can reduce group g real wages and employment if (i) workers cannot easily reallocate (θ_{gg} is high) and (ii) the cost-saving gains from automation π_{ng} are small—as in “so-so” automation technologies. Both conditions are necessary. In the limit where workers can reallocate by taking small wage cuts, $\theta_{gg} = \theta_{jg}$ for all j , and all wages change by the same amount. Solow’s dual in (7) implies that all real wages increase. Conversely, if the cost-saving gains from automation π_{ng} are large, the productivity effect dominates the displacement effect for all workers, increasing wages.

Taking stock: The results in this section explain how automation can contribute to recent trends. They show that automation is a plausible driver of the observed labor share

¹⁶Structural transformation, trade in final goods, and reallocation across firms affect the demand for skills through this channel (see Buera et al., 2021, for work on the role of sectoral shifts).

decline in the aggregate and especially in manufacturing. They also show that automation can explain the changing occupational structure in the US and Europe, featuring a pronounced decline in the employment in routine middle-pay occupations, as these occupations contain tasks that can be more easily automated given recent technological advances. Finally, they show that the shifts in the US wage structure can be due to the automation of tasks previously performed by non-college workers that brought small productivity gains and left them with limited opportunities to reallocate. The formulas also emphasize the role of the displacement effects from automation and explain why this key feature is necessary to account for the observed trends.

I conclude this section by discussing the different implications of automation and other forms of technological progress, and by comparing the aggregate effects of automation on wages and employment to firm-level estimates found in the empirical literature. These distinctions are important for interpreting recent empirical work on the impact of automation.

Automation vs. Other Technologies

One important feature of the task model is that it recognizes that technology is multifaceted and can affect production via distinct margins. In particular, the task model distinguishes automation from other forms of technological progress that do not create displacement effects, including the creation of new tasks and products and advances in the productivity of capital at the intensive margin.¹⁷

The creation of new tasks and products increases output by providing workers with additional (or better) productive opportunities (as in [Acemoglu and Restrepo, 2018](#); [Hémous and Olsen, 2022](#)). One can model new product creation as an increase in \mathcal{N} , which raises output via a love-for-variety effect and shifts the demand for labor in favor of skills used in new products. New labor-intensive products can raise the labor share and benefit all workers, offsetting the displacement effects from automation.

Advances in capital at the intensive margin increase the productivity (or reduce the cost) of task-specific capital in use (as in [Acemoglu and Restrepo, 2019](#); [Jones and Liu, 2022](#)). For example, firms may replace older vintages of capital with newer ones performing the same task more effectively (think of firms replacing a steam engine with an electric motor

¹⁷In principle, one can view these processes as separate from automation. In practice, broader technological developments may affect production through various margins. For example, the arrival of computers led to the development of automated software systems used for automation, but computers also facilitated the introduction of new products and services. Even in these cases, it is helpful to distinguish conceptually and empirically between automation and other margins, as these have different implications.

or upgrading cranes and conveyors to more rapid ones). Or firms may use existing machines more intensively when their price drops due to advances in capital production.

None of these forms of investment displace workers from their tasks, distinguishing them from automation. Equation (4) illustrates the distinction. Intensive-margin advances in capital are isomorphic to an increase in k_n —these are equivalent to having firms operate with more capital holding their CES weights constant. This differs from automation, which maps to a direct shift in the CES weights as firms automate tasks at the extensive margin.

Unlike automation, intensive-margin capital advances do not produce any direct shift in adopting firms’ occupational structure and only affect firm and industry labor shares via task prices, reducing it in the relevant case with $\lambda < 1$. Intensive margin advances can also benefit all workers equally—even if they cannot reallocate across tasks. For example, in the case analyzed above with no product shifters and a common row sum for Θ , intensive-margin advances in capital increase all group wages and employment by the same amount.

The distinction between automation and other technologies is relevant for interpreting historical patterns and recent trends. For example, [Acemoglu and Restrepo \(2018\)](#) and [Jones and Liu \(2022\)](#) show that ongoing automation in the past can be consistent with the historical balanced growth experience if accompanied by the creation of new labor-intensive tasks and/or intensive-margin advances in machinery. Deviations from balanced growth result from automation outpacing these countervailing forms of technological progress.

The distinction is also relevant for empirical work, highlighting the importance of separating automation from other forms of technological progress to identify their distinct implications for firm and worker outcomes.

Reallocation and Firm and Industry Employment.

A growing body of work estimates the impact of automation on firm sales and employment. This section derives the direct effects of automation on firm sales and employment and discusses their interpretation.

Consider a scenario where n denotes firms producing differentiated goods y_n . The direct effect of automation on firm n sales compared to non-adopters is

$$(10) \quad \delta \ln(p_n \cdot y_n) = (\epsilon_n - 1) \cdot d \ln t f p_n,$$

and the direct effect on firm n employment compared to non-adopters is

$$(11) \quad \delta \ln \ell_{ng} = \underbrace{(\epsilon_n - \lambda) \cdot d \ln tfp_n}_{\text{Scale vs. substitution effects}} - \underbrace{d \ln \Gamma_{ng}^d}_{\text{displacement effect}}$$

Here, $\epsilon_n > 1$ is the demand elasticity faced by firm n .

The first term in (11) captures scale and substitution effects. Scale effects result from sales reallocation from non-adopters to adopting firms holding aggregates constant (governed by the demand elasticity ϵ_n). The substitution effect results from firms substituting away from labor-intensive tasks that have not been automated (governed by λ). Reallocation dominates when $\epsilon_n > \lambda$, which is a common empirical configuration. The second term in (11) captures the negative contribution of displacement effects to firm-level employment. This term is unique to automation. For example, advances in capital at the intensive margin only affect firm employment via scale and substitution effects and increase it if $\epsilon_n > \lambda$.

Equation (11) shows that cross-sectional comparisons of firms provide information on whether reallocation dominates the displacement and substitution effects. The net impact on firm-level employment depends on the strength of reallocation across firms, governed by ϵ_n . For low demand elasticities, automation can simultaneously increase firm sales and reduce employment (and this holds even for $\epsilon_n > \lambda$). With highly elastic demand, adopters will see employment growth relative to non-adopters, although at a slower rate than sales.

Though relevant for some questions, firm-level estimates are not informative of the aggregate consequences of automation for wages and employment of displaced workers (shown in equation 7). First, reallocation across firms (the product-mix term in 7) has no clear aggregate implications. It can be neutral for aggregates if $\omega_g^n = s_y^n$ for all g , so that firms employ inputs in equal proportions. And yet, it but would continue to show up positively in firm-level effects. Second, as discussed above, the key determinants of the general equilibrium effects of automation on employment and wages are (i) how strong the productivity effect $d \ln y$ is relative to displacement effects and (ii) how easily workers reallocate. Cross-sectional firm estimates entirely miss these forces.¹⁸

The same formulas and limitations apply to industry-level effects. The main difference is that when n denotes industries, the elasticity ϵ_n can be below one if goods are complements. In this case, industry-level estimates capture the negative reallocation away from adopting

¹⁸This discussion also clarifies that the emphasis on demand elasticities ϵ_n as a key determinant of the impact of automation (and technology more broadly) on aggregates is misplaced. Demand elasticities determine the strength of reallocation across firms and industries, which has no clear aggregate implications.

industries, which has no clear aggregate implications either and cannot be taken as evidence of a negative impact of automation on displaced workers' wages and employment.

2 EMPIRICAL WORK ON AUTOMATION

This section reviews the empirical literature exploring the implications of different automation technologies for factor shares, occupational structure, exposed workers, and firm and industry-level outcomes.

Automation and the labor share

The prediction that automation reduces the labor share in value added and increases sales per worker has been documented in various contexts by studies exploring the impact of automation technologies on industries and firms. Most of this literature has studied the impact of industrial robots, which provide a clear example of an automation technology. Recent contributions have extended these findings to other technologies.

[Acemoglu and Restrepo \(2020\)](#) studied the introduction of industrial robots across US industries in the 1990s and 2000s. They show that the adoption of industrial robots in the US was driven by technological developments abroad (in Japan, Germany, and other European countries). They document that US industries benefiting from advances in industrial robotics increased value added and reduced employment. This resulted in higher output per worker and a lower labor share. Moving one industry from zero robots to 10 robots per thousand workers (the level in metalworking industries) increases value added by 12.5%, reduces employment by 6.25%, and lowers its labor share by 5 pp.

[Graetz and Michaels \(2018\)](#) explored this link for a broader set of European countries with higher levels of robot adoption than the US. They show that industries intensive in tasks that can be automated via industrial robots (picking, reaching, and others) adopted more robots and saw a greater increase in value added per worker from 1993 to 2007. They estimate that moving from the lowest (zero robots) to the highest level of adoption (the level in car manufacturing) increases value added per worker in an industry by 66–100 log points and raises average wages by 10 log points. This implies a 56-90 log point decrease in labor shares. Their appendix tables document that these results are accompanied by a 50–60 log point increase in value added and a 20–50 log point decrease in employment (though their employment estimates are imprecise and cannot rule out zero or large negative effects).

These findings align with the growing literature exploiting the adoption of industrial robots across firms discussed below, which finds that this technology is associated with declining firm labor shares.

Recent work has turned to a broader range of automation technologies.

[Boustan et al. \(2022\)](#) studied the introduction of computer-numerically-controlled (CNC) machinery in the US, which preceded the arrival of industrial robots. CNC machinery (including lathes, milling machines, and others) automated the role of semi-skilled machine operators in metalworking industries. As in the case of industrial robots, the adoption of CNC machinery in the 1970s and 1980s was driven by technological developments abroad (most notably in Japan and Germany). Manufacturing industries benefiting from these developments saw declining labor shares. A 10 pp increase in the share of imported CNC machinery in an industry (among all machine tools) is associated with a 1.6 pp decline in its labor share and a 20% increase in output per worker during 1960–2010. The increase in output per worker results from a 24% increase in output and a 4% decrease in employment, though their employment estimates cannot rule out zero effects.

[Acemoglu and Restrepo \(2022\)](#) considered the role of dedicated machinery (including CNC machines and other automatic machinery with specific functions, such as self-checkouts and ATMs) and specialized software systems (including custom software developed for inventory, customer, and human resource management). Using BLS detailed asset tables for 49 US industries, they compute the change in the share of dedicated machinery services and specialized software services for 1987–2016. They find that 50% of the labor share decline across US industries during this period can be explained by the increased use of dedicated machinery, specialized software services, and industrial robots. This relationship is robust to controlling for changes in markups, rising sales concentration, and declining unionization rates across industries.

[Kogan et al. \(2021\)](#) and [Dechezleprêtre et al. \(2023\)](#) provide additional evidence on the impact of automation technologies on the labor share using text-analysis techniques. [Kogan et al. \(2021\)](#) create a measure of similarity between the capabilities of *breakthrough innovations* (based on the text of highly-cited patents' descriptions) and tasks in an occupation (from ONET). Their analysis focuses on high-similarity breakthrough innovations, which are presumably the ones capable of substituting for worker tasks. Using NBER-CES manufacturing data for 1958–2018, they show that a one-standard-deviation increase in the number of high-similarity patents in an industry is associated with a 2.8% increase in labor productivity and a 1.25% decline in its labor share over five year periods. [Dechezleprêtre](#)

et al. (2023) use text analysis to classify patents as automation if they substitute for labor in some tasks. They find that a 1 pp increase in the share of automation patents in an industry is associated with a 1.3 pp decline in its labor share (for reference, their share of automation patents increased by 10 pp from 1980 to 2015 in most countries).

The idea that automation contributed to the labor share decline also aligns with broader sectoral patterns.¹⁹ In the US, most of the labor share decline concentrates in manufacturing—the sector with the greatest adoption of automation technologies (see Acemoglu et al., 2022)—and in equipment and software-intensive sub-industries (see Hubmer, 2023).²⁰

Automation and the changing occupational structure

The available evidence supports the view that advances in automation contributed to the shift away from middle-pay occupations, including routine cognitive and manual jobs, as these were highly exposed to advances in automation technologies over the last 50 years.

Autor et al. (2003) were the first to emphasize this possibility. They argue that computer capital substitutes for workers in routine tasks because these follow codifiable rules. Using US data for 1960–2000, they document that the decline in routine tasks concentrated in rapidly computerizing industries and that these shifts accelerated in the 1970s as computer prices declined. Since then, a vast literature has documented a decline in routine jobs in various countries (see Acemoglu and Autor, 2011; Goos et al., 2014). Dechezleprêtre et al. (2023) show that the decline in routine occupations is also more pronounced in industries with the greatest increase in automation patents.

Recent contributions have moved beyond the classification of occupations into routine jobs and have used text analysis to measure the extent to which new technologies affect occupations. These studies produce indices of *occupational exposure* to technological advances, defined by the similarity between tasks in an occupation (obtained from text descriptions in ONET or similar sources) and the capabilities of new technologies described in patent documents. The robust finding in these papers is that exposed occupations have seen a sizable employment decline over time, contributing to the changing US occupational

¹⁹An alternative view is that the labor share decline is due to rising markups. This runs counter to the fact that most of the decline in the US labor share is in manufacturing—the sector with the smallest increase in sales concentration and estimated increase in markups (see Hubmer and Restrepo, 2021).

²⁰One challenge to the view that automation contributed to the labor share decline is that the labor share of a typical firm has not decreased. Autor et al. (2020) and Kehrig and Vincent (2021) document that, while the aggregate labor share in the US declined, the labor share of the median firm and an unweighted mean of labor share changes across firms remained unchanged since the 1980s. Hubmer and Restrepo (2021) show that a model with a fixed cost of automation per task can reproduce these facts.

structure. A reasonable interpretation of these findings is that exposure indices capture technology’s capacity to substitute for worker tasks in an occupation.

[Webb \(2020\)](#) used this approach to create indices of occupational exposure to robotics, software, and artificial intelligence. His indices show that exposure to robotics is high for middle and low-pay blue-collar occupations, and exposure to software is high for middle-pay white-collar occupations. He estimates that moving from the median (technicians) to the highest (machine feeders) percentiles of robot exposure is associated with a decline in employment of 20% and wages of 15% during 1980-2010. Similarly, moving from the median (economists) to the highest (power-plant operators) percentiles of software exposure is associated with a decline in employment of 7–15% and wages of 2–6.5%.

[Kogan et al. \(2021\)](#)’s work listed above measures occupational exposure to breakthrough innovations for 1900–2000. They find that middle-pay occupations were more exposed to breakthrough innovations, with routine-manual jobs more exposed early on and routine-cognitive jobs more exposed after 1980. Using Census data for 1910-2010 and CPS data for 1983–2010, they show that a one standard deviation increase in exposure is associated with a 20% decrease in employment and a 4% decrease in wages in the next 20 years.

[Autor et al. \(2022\)](#) develop text-based measures of occupational exposure to automation technologies and augmenting technologies. Their automation measure is computed from the text similarity between patents and tasks in a job. Their augmentation measure is computed from the text similarity between patents and job titles in an occupation, obtained from the *Census Alphabetical Index of Occupations and Industries*. The argument is that job titles capture the services rendered in an occupation (and not the tasks involved in rendering these services). Related patents presumably capture innovations that help workers render these services. Using US data for 1940–1980 and 1980–2018, they show that the share of employment and wage payments expanded in occupations exposed to augmenting innovations and contracted for those exposed to automation innovations. A one standard deviation increase in automation patents is associated with an 8–16% decline in employment for exposed occupations.

The evidence from these studies aligns with work studying how the adoption of robots or CNC machinery impacted the employment composition of industries and firms. For example, [Acemoglu and Restrepo \(2020\)](#) and [Boustan et al. \(2022\)](#) document that US manufacturing industries benefiting from advances in industrial robotics and CNC machinery saw a decrease in the share of workers employed in blue-collar routine occupations. The evidence for firms discussed below supports a similar conclusion.

Automation and group-level outcomes

The task model predicts that groups displaced from their tasks by automation will experience a relative decline in wages and employment, especially if they cannot reallocate. The literature has explored this implication by tracing the impact of automation on individually exposed workers, workers in exposed regions, and workers in exposed skill groups.

Evidence using panel data to trace the outcomes of individual workers exposed to automation includes work by Cortes (2016), Kogan et al. (2021), Bessen et al. (2023), and work from a unique historical context by Feigenbaum and Gross (2020). These designs estimate the incidence of displacement effects on individual workers directly exposed to automation and study their adjustment. One can think of these settings in the task model by letting g index groups of incumbent workers in an occupation, industry, or firm that adopts automation technologies.

Cortes (2016) uses data from the Panel Study of Income Dynamics to explore how workers in routine occupations adjusted to the automation of these jobs. He documents that although some incumbents reallocated, exposed workers suffered a sizable income loss. On average, workers who held routine jobs in 1980 saw a 17% income decline over the next 20 years relative to non-exposed workers.

Kogan et al. (2021) use data from the US Social Security Administration to trace workers in occupations exposed to labor-replacing technologies, according to their text-based measure. Workers in exposed occupations experienced a 2.3% income decline (relative to non-exposed workers) ten years after a one-standard-deviation improvement in labor-replacing technology.

Bessen et al. (2023) study the impacts of firm investments in automation technologies on incumbent workers using Dutch employer-employee matched data for 2000–2016. They use a unique survey reporting firm expenditures on third-party automation technology providers. These include payments made to integrators—companies offering engineering and software solutions for various automation technologies. They document that investments in integration are lumpy and take place in spikes. These spikes provide a compelling proxy for the adoption of automation technologies at the extensive margin since significant reorganizations of production require assistance from integrators. In the five years following an investment spike in automation services, incumbent workers experience an increase in separation rates and a cumulative labor income loss totaling 10% of their annual income relative to non-exposed workers.

Feigenbaum and Gross (2020) study the effects of mechanizing telephone operation in the US between 1920 and 1940 on incumbent workers and new cohorts of young women. They identify women employed as telephone operators before this shock and trace their labor-market outcomes over time by linking decennial Censuses. They find that after a city switches to mechanical operation, the number of young women employed as operators immediately declined by 50–80%. Incumbent (female) workers experience a subsequent decline in employment of 8 pp, providing evidence of the negative consequences of the displacement effect on exposed workers. Even though some managed to reallocate to clerical jobs in other industries, they were generally forced into lower-paying occupations. New cohorts of women reallocated more swiftly and experienced no decline in employment, though some might have taken lower-paying jobs.

These studies support the view that workers whose tasks and jobs are automated experience worse labor market outcomes subsequently relative to other workers, as they are left with fewer employment opportunities. Workers’ ability to reallocate mitigates these adverse impacts to some degree, though this mechanism is most relevant for young workers in new cohorts. However, these designs miss broader equilibrium impacts on non-incumbent workers with similar skills operating through task competition with directly displaced workers.

A second set of studies estimates the impact of automation on local labor markets, such as US commuting zones. One can think of this setting in the task model by letting g index the group of workers in a local labor market.²¹ These studies cannot tell us what happens to individual workers who are displaced but are informative of the broader equilibrium effects on regional outcomes.

Acemoglu and Restrepo (2020) use this approach to explore the implications of advances in robotics on exposed US regions. They find that from 1990 to 2007, US commuting zones that specialized in industries experiencing significant advances in industrial robots saw a relative decline in employment and wages. Advances leading to the adoption of one industrial robot per thousand workers in a commuting zone reduce wages by 0.8 percent and its employment-to-population ratio by 0.4 pp relative to other regions, with 0.15–0.2 pp of the employment decline coming from manufacturing. This implies a reduction of 2–3 manufacturing jobs per robot in exposed regions. This evidence is consistent with a world where workers in exposed regions suffer the displacement effects from automation while

²¹With this interpretation, the propagation matrix captures the effects of migration across regions, as in work by **Borusyak et al. (2022)**. Due to low migration flows (at least in the US), the propagation matrix will be close to diagonal, which implies that exposed commuting zones bear most of the incidence of automation and other shocks to their labor demand.

productivity gains are shared nationally.

[Dauth et al. \(2021\)](#) extend this approach to Germany. They estimate that technological advances leading to an extra robot in a region reduce its manufacturing employment by two jobs—a magnitude comparable to [Acemoglu and Restrepo \(2020\)](#). However, they find no adverse effects on total regional employment, as young workers entering the labor market find jobs in the expanding business services sector. One plausible interpretation is that the displacement effects from robot adoption in Germany and the US are of a similar magnitude, but in Germany, the productivity gains from automation are higher and benefit exposed labor markets the most. This could be because advances in robotics lead to the expansion of integrators and robot producers in exposed regions.

[Boustan et al. \(2022\)](#) study the local-labor market effects of advances in CNC machinery from 1970 to 2000 in the US. A one standard deviation increase in exposure to CNC machinery (instrumented by advances abroad in Germany and Japan) for a commuting zone reduces the share of the population employed in metal-working manufacturing industries by 3.5 pp but has no adverse effect on total manufacturing employment. This, too, points to a displacement effect in metal-working manufacturing that is offset by productivity gains benefiting other local manufacturing industries.

[Mann and Puttmann \(2023\)](#) study the local labor market effects of advances in automation identified from patent data. They find that US commuting zones that specialize in industries with a higher rate of automation patenting saw gains in employment. [Coelli et al. \(2023\)](#) revisit this conclusion using an updated classification of automation patents and report a negative employment effect in exposed industries and US commuting zones over the 1980-2010 period. The different findings reflect their measurement of automation patents. [Mann and Puttmann \(2023\)](#) classify a patent as automation if it refers to a technology capable of performing functions independently. This definition captures mainly innovations in computers and communications that do not automate tasks performed by workers, with 90% of patents in computers and communications labeled as automation. [Coelli et al. \(2023\)](#) use the classification from [Dechezleprêtre et al. \(2023\)](#), which codes a patent as automation if it describes an innovation capable of substituting for current worker tasks, matching the definition in the task model.

The third set of studies estimates the impact of automation on entire skill groups, defined by workers' demographic and educational characteristics. While these designs cannot tell us what happens to individual workers who are displaced, they are informative of the broader equilibrium effects on group-level outcomes at the national level.

[Acemoglu and Autor \(2011\)](#) pioneered this approach. Their work shows that demographic groups of workers specialized in routine occupations experienced a subsequent decline in wages from 1960 to 2010.

[Acemoglu and Restrepo \(2022\)](#) further developed this approach to estimate the impact of automation on group-level wages and employment, using data for 500 groups of US workers. These groups are defined by gender, age, education, race, and birthplace, though they also provide robustness checks expanding the definition of their groups to account for region of residence. Their approach relies on measuring the direct task displacement experienced by each group from 1980 to 2016 (i.e., $\sum_n \omega_n^g \cdot d \ln \Gamma_{ng}^d$ in the model) and then regressing the change in group-level outcomes during this period on this measure. Their measures of direct task displacement are associated with large reductions in group wages and employment. A 10% increase in direct task displacement (i.e., a reduction in the task share of group g by 10% due to automation) is associated with a 15% reduction in group relative wages and a 4.4 pp decline in their employment rate during 1980–2016. This effect is robust to controlling for educational dummies and industry shifters. Moreover, their task displacement measure explains 50% of the observed change in group wages in this period.

Automation and firm-level outcomes

A growing literature has turned to firm-level data to study the implications of automation technologies. The robust finding is that firms adopting automation technologies increase their sales and employment but see a reduction in their labor share and a change in their workforce composition.

[Acemoglu et al. \(2020a\)](#) use data for French manufacturing firms for 2011–2014 and document that robot users expanded value added by 20%, employment by 10%, and saw no change in average wages. This results in a 10% (or 4.3 pp) reduction in their labor share. Other papers on robot adoption find similar estimates. Using data for Spanish manufacturing firms, [Koch et al. \(2021\)](#) find that robot adoption is associated with a 25% increase in sales, a 10% increase in employment, and a 6.5 pp decline in the labor share of value added. Using data for Danish manufacturing firms, [Humlum \(2020\)](#) finds that robot adoption events are associated with a 20% increase in sales, a 10% increase in employment, and a 10% decline in the labor share. Using data for Dutch manufacturing firms, [Acemoglu et al. \(2023b\)](#) report similar estimates. Recent work by [Bonfiglioli et al. \(2020\)](#) using French data also finds a 13% increase in labor productivity associated with robot adoption, driven

by a 23% increase in sales and a 10% increase in employment.²²

These papers also document a shift in the composition of employment away from production and routine-manual jobs and an increase in the demand for skilled labor. The work by Humlum (2020) for Danish firms stands out for having high-quality data on employment by detailed occupation at the firm level. He finds that in the five years following the adoption of industrial robots, firms see a 20% reduction in the share of wage payments to workers in production jobs (such as assembly or welding).

Other studies support the idea that a broader set of automation technologies, including CNC machinery or specialized software, reduce firms' labor shares and change their workforce composition, though the evidence is not as conclusive as for robots.

Cheng et al. (2021) use data for China for 2015–2018 and exploit city-level variation in a government program subsidizing the adoption of industrial robots and CNC machinery for identification. Their estimates imply that a 10% subsidy for investments in robots and CNC is associated with a 2–3.5 pp decline in firms' labor shares.

Dinlersoz and Wolf (2023) use the *Survey of Manufacturing Technologies (SMT)* from 1991 and show that in the cross section of US manufacturing plants, those using automation technologies have lower labor shares and employ a lower share of production workers. Acemoglu et al. (2022) use data from the *Annual Business Survey (ABS)* for 2016–2018 and document that US firms using robotics, specialized software, and dedicated equipment in 2016–2018 had lower labor shares and higher sales per worker. The evidence in these two papers is descriptive, as it relies on cross-sectional comparisons. Nonetheless, it aligns with the fact that a significant share of firms in the SMT and ABS report adopting these technologies to automate workers' tasks. For example, the ABS data shows that 30.4% of US workers and 50% of US manufacturing workers in 2016–2018 were employed at firms using some advanced technology for automation.

A notable aspect of the firm-level evidence is that the majority of papers estimate an expansion of employment at firms adopting automation technologies. This is what we should expect if firms operate in highly competitive environments with very elastic demands, causing adopting firms to expand at the expense of their competitors—as documented in Acemoglu et al. (2020a) and Koch et al. (2021). The theory clarifies that this reallocation across firms has no straightforward aggregate consequences and that because

²²One limitation of these studies is that adopters may be on a more steep growth trajectory than competing firms (as shown in Acemoglu et al., 2023a, for the US) or that firms may time their adoption to periods of high demand.

of this, firm-level employment estimates do not tell us much about the aggregate impact of automation on the wages and employment of displaced workers.

Automation vs. other investment margins

The task model draws a clear distinction between the implications of firms adopting new automation technologies and the effects of investments in capital at the intensive margin or in technologies associated with new goods and products. This distinction is relevant for interpreting recent work by [Curtis et al. \(2021\)](#), [Hirvonen et al. \(2022\)](#), and [Aghion et al. \(2023\)](#), who report no effects of increased firm investment in response to tax benefits or lower equipment prices on firms' labor shares and skill composition.

[Curtis et al. \(2021\)](#) estimate the impact of accelerated depreciation policies on investment and firm-level employment in the US from 2001 to 2011. They find that firms that benefited from the policy increased capital investment and employment relative to a control group but saw no decline in their labor share nor increases in their demand for skills.

[Hirvonen et al. \(2022\)](#) study a program that subsidized investments in dedicated equipment in Finland. They report an expansion in firm sales and employment but no changes in firm labor shares or workforce composition.

[Aghion et al. \(2023\)](#) use French manufacturing data to study the implications of investments in imported automation equipment. Their main design exploits improvements in foreign suppliers' productivity as an exogenous shock increasing imports of automation equipment. Using supplier shocks as instruments, they estimate an expansion in firm and industry employment but find no effects on firms' labor shares or the share of wages paid to production and non-college workers.

One way of understanding these findings through the lens of the task model, is as reinforcing the point that not all investments in manufacturing technologies generate displacement effects. More precisely, these results align with the view that investments in capital at the intensive margin or in machinery for new products generate no displacement effects and are, therefore, qualitatively different from investments in automation technologies at the extensive margin. For example, the accelerated depreciation scheme in [Curtis et al. \(2021\)](#) can induce firms to use more capital at the intensive margin, especially if perceived as a temporary benefit. The subsidy program studied in [Hirvonen et al. \(2022\)](#) worked by getting firms to purchase machinery for new products. [Aghion et al. \(2023\)](#)'s approach using supplier shocks exploits variation for firms that have already imported

automation equipment from a supplier, isolating the effect of further investments at the intensive margin.²³ These forms of investment create no displacement effects, explaining why these studies find no changes in firms’ labor shares, demand for skills, and occupational structure and why the results here differ from previous findings.²⁴

3 AREAS FOR FUTURE WORK

Improving the measurement of automation

The empirical evidence provides qualitative support for the implications of automation in task models. However, most of the evidence is reduced form and relies on indirect proxies for automation without clear magnitudes, such as ordinal indices of occupational exposure. The few instances where we have measures of actual investment in automation technologies are confined to specific technologies such as industrial robots. For this reason, existing empirical work does not yet provide clear-cut quantitative answers to questions such as “What share of worker tasks has been automated in the last ten years?” or “Has the rate of automation accelerated since the 1980s?” To address these questions, the literature needs to develop better measures of automation and quantitative strategies to go beyond reduced-form analyses.

Measuring automation The objective is to develop a systematic approach to identify automation technologies, measure expenditures in these technologies and their direct task displacement, and quantify the productivity gains associated with their deployment.

One approach would proceed as follows:

- Identify technological developments that lead to the automation of tasks performed

²³Aghion et al. (2023) offer a different interpretation. They see the increase in employment as evidence of a strong scale effect that dominates the displacement effect, in part because French firms compete in international markets. Either way, the findings in these papers have clear policy relevance: they show that blanket taxes on capital and equipment imports have limited distributional benefits and distort investments that do not generate displacement effects (see Acemoglu et al., 2020b; Donald, 2022, for work exploring the implications of policy tools that can target the extensive margin of automation).

²⁴Previous work by Acemoglu and Restrepo (2020); Feigenbaum and Gross (2020); Kogan et al. (2021); Autor et al. (2022); Boustan et al. (2022) isolates the extensive margin by identifying the arrival of new automation technologies and estimating the effects of their initial adoption and deployment across firms, industries, occupations, or regions. Firm-level studies on the implications of robots isolate the extensive margin by identifying the time at which firms became robotized, and tracing firm outcomes following this event. The exception is Acemoglu et al. (2020a) who use data for 2011-2015 and cannot distinguish between firms adopting robots for the first time (extensive margin) or purchasing additional robots over time without automating more processes (intensive margin).

by workers, either by using patent data (as in [Kogan et al., 2021](#); [Autor et al., 2022](#); [Dechezleprêtre et al., 2023](#)) or by relying on the engineering literature and historical accounts of what specific technologies do (as in [Acemoglu and Restrepo, 2020](#); [Boustan et al., 2022](#); [Feigenbaum and Gross, 2020](#)).

- Confirm that we have identified an automation technology by correlating its initial adoption across firms or industries with a declining labor share and rising labor productivity, as many papers do via reduced-form analyses.
- Measure investments (both in quantities and value) at using firms and industries in this technology. This key missing step can be accomplished by using customs data or detailed firms’ balance sheets to measure expenditures in equipment or software embodying the technology. This can also be accomplished by using firm-to-firm trade data to measure sales by customer for technology providers, integrators, and third-party providers of automation solutions or by directly surveying providers.²⁵

The resulting measures of adoption and investment can then be used for empirical work and quantitative analyses. For example, researchers can trace the share of expenditure in automation technology in costs over time and across firms and industries. Long time series would allow researchers to estimate the total displacement and productivity gains from this technology, starting at its insertion and tracing its impact as firms adopt it for the first time. This is important for distinguishing investments at the extensive margin early on (firms adopting the new technology and re-organizing their production process) from subsequent investments at the intensive margin (firms upgrading the initial machinery and software later in time, adding to its productivity effect).

An alternative approach involves using *operations data*. These data provide establishment-level information, breaking the production of goods into tasks and describing how each is completed (workers and tools involved, machinery if automated, power sources, time requirements, and so on). This provides all the information needed to measure task shares and identify automation advances. Examples include the BLS “Hand and Machine Labor Study” from the mid-1890s (analyzed in [Atack et al., 2019](#)), and hand-collected data for establishments producing cars and semiconductors (in [Ales et al., 2023](#)). These data are rare and available only for some sectors and products. However, firms are increasingly collecting operations data, which could make this approach feasible at scale in coming years.

²⁵This is the approach followed by the *International Federation of Robotics* (IFR). The IFR surveys the main world suppliers of industrial robots (a handful of firms) and creates statistics on robot installations by year, country, and industry from their sales-by-customer data. One reason why this approach is feasible is that technology suppliers are highly concentrated.

Going beyond reduced-form estimates One fruitful strategy is to combine proxies for automation with accounting data to estimate the direct task displacement $\{d \ln \Gamma_{ng}^d\}$ and cost-saving gains $\{\pi_{ng}\}$ associated with the adoption of robots and other automation technologies. As explained in the theory section, these estimates provide sufficient statistics for the capabilities and productivity of automation technologies.

These approaches can use firm data, as in [Humlum \(2020\)](#), or industry data, as in [Acemoglu and Restrepo \(2022\)](#) to infer the change in task shares and cost-saving gains brought by the introduction of automation technologies from firm or industry-level impacts.

[Humlum \(2020\)](#) assumes that firms operate a CES production function that combines workers in different occupations and allows the adoption of industrial robotics to shift CES weights and increase TFP.²⁶ His approach infers the shifts in weights across occupations—corresponding to $\{d \ln \Gamma_{ng}^d\}$ in the model—and the productivity gains brought by industrial robots from the behavior of sales and wage shares in adopting Danish firms. These estimates can then be used to compute the impact of observed robot adoption and conduct counterfactuals.

As a second example of how firm data can be used more effectively, consider the 20% sales expansion and 10% labor share reduction found for European firms adopting industrial robotics. Assume that industrial robots substitute for non-college workers, that these workers represent 70% of firm wages, that labor accounts for 25% of gross costs, and that firms face a demand elasticity of 10. These estimates then imply that robot adoption reduces the task share of non-college workers in adopting firms by $d \ln \Gamma_{ng}^d = 16.8\%$ (to match the labor share decline in [5](#)) and costs at automated tasks by $\pi_{ng} = 70\%$ (to match the sales expansion in [10](#)). These estimates can be plugged in equations [\(7\)](#), [\(8\)](#), and [\(9\)](#) to compute aggregate effects and conduct counterfactuals.

[Acemoglu and Restrepo \(2022\)](#) propose a method for estimating the direct task displacement experienced by groups of US workers. Their method infers the direct task displacement experienced by group g in industry n , $d \ln \Gamma_{ng}^d$, from the percent decline in the industry labor share explained by automation proxies. This is computed as the predicted value from the cross-industry regression described above, explaining the labor share decline in an industry as a function of its use of dedicated machinery, specialized software services, and industrial robots. This is then apportioned across groups in proportion to their revealed comparative advantage in routine jobs in that industry.

²⁶This is a specification of [\(4\)](#) where groups are defined by occupations, and the propagation matrix is the identity. One challenge when interpreting groups as occupations is that workers change occupations over time. [Humlum \(2020\)](#) addresses this by modeling occupational choice as in [Traiberman \(2019\)](#).

Their measure implies that workers at the bottom and middle of the wage distribution lost 20–30% of their tasks since 1980 to automation, while workers with a post-college degree experienced almost zero displacement. Their paper also develops a methodology for estimating the propagation matrix, which allows them to compute the effects of automation on the US wage structure using equations (7), (8), and (9). They find that 50% of the changes in the US wage structure between educational and demographic groups since 1980 can be explained by the uneven incidence of the displacement effects from automation.

Future work could build on these approaches by estimating or measuring the direct task displacement and productivity gains associated with different automation technologies, improving the modeling of the endogenous reassignment of tasks and the estimation of the propagation matrix, extending these methods to account for larger shocks, or proposing flexible parameterizations of the task model suitable for a full structural analysis.

Artificial Intelligence

An important area for future work is to explore the extent to which insights from existing studies can inform predictions about the influence of Artificial Intelligence (AI) on labor markets and the workforce. One fundamental question that remains unanswered is whether it is appropriate to model narrow AI systems as automation technologies that could perfectly replace human labor in specific tasks. Alternatively, it could be argued that narrow AI systems are not tools for automating tasks but are instead augmenting the abilities of human workers who use them, widening the supply of skills in the economy.

Although research in this field is still nascent, emerging patterns offer some grounds for informed conjecture. For instance, recent studies indicate that Large Language Models (LLMs) may be capable of substituting human labor in writing tasks, lending credence to the notion that AI may automate a wider array of tasks in the short term (see [Noy and Zhang, 2023](#)). However, differently from recent waves of automation, the types of tasks that AI (and in particular LLMs) can perform are more evenly distributed across workers in many occupations of different skill levels (see [Webb, 2020](#); [Eloundou et al., 2023](#)). If anything, high-pay workers appear more exposed to any potential displacement effects. According to the task model, the impact of AI-driven automation may not result in the same negative distributional consequences observed in previous automation waves. This conclusion is of course contingent on the specific tasks that AI will ultimately be able to automate as these systems gain more advanced capabilities and applications.

REFERENCES

- ACEMOGLU, DARON, GARY ANDERSON, DAVID BEEDE, CATHERINE BUFFINGTON, ERIC CHILDRESS, EMIN DINLERSOZ, LUCIA FOSTER, NATHAN GOLDSCHLAG, JOHN HALTIWANGER, ZACHARY KROFF, PASCUAL RESTREPO, AND NIKOLAS ZOLAS (2022): *Automation and the Workforce: A Firm-Level View from the 2019 Annual Business Survey*, U. Chicago Press.
- (2023a): “Advanced Technology Adoption: Selection or Causal Effects?” *AEA Papers and Proceedings*, 113, 210–14.
- ACEMOGLU, DARON AND DAVID AUTOR (2011): *Skills, Tasks and Technologies: Implications for Employment and Earnings*, Elsevier, vol. 4 of *Handbook of Labor Economics*, chap. 12, 1043–1171.
- ACEMOGLU, DARON, HANS R. A KOSTER, AND CEREN OZGEN (2023b): “Robots and Workers: Evidence from the Netherlands,” Working Paper 31009, National Bureau of Economic Research.
- ACEMOGLU, DARON, CLAIRE LELARGE, AND PASCUAL RESTREPO (2020a): “Competing with Robots: Firm-Level Evidence from France,” in *AEA Papers and Proceedings*, vol. 110, 383–88.
- ACEMOGLU, DARON, ANDREA MANERA, AND PASCUAL RESTREPO (2020b): “Does the US Tax Code Favor Automation?” *Brookings Papers on Economic Activity*, 2020, 231–300.
- ACEMOGLU, DARON AND PASCUAL RESTREPO (2018): “The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment,” *American Economic Review*, 108, 1488–1542.
- (2019): “Automation and New Tasks: How Technology Displaces and Reinstates Labor,” *Journal of Economic Perspectives*, 33, 3–30.
- (2020): “Robots and Jobs: Evidence from US Labor Markets,” *Journal of Political Economy*, 128, 2188–2244.
- (2022): “Tasks, automation, and the rise in US wage inequality,” *Econometrica*, 90, 1973–2016.
- (2023): “Automation and Rent Dissipation:,” Tech. rep., Mimeo, Boston University.
- AGHION, PHILIPPE, CELINE ANTONIN, SIMON BUNEL, AND XAVIER JARAVEL (2023): “Modern Manufacturing Capital, Labor Demand, and Product Market Dynamics: Evidence from France,” Tech. rep., CEPR Discussion Papers No. 1910.
- AGHION, PHILIPPE, BENJAMIN F JONES, AND CHARLES I JONES (2018): “Artificial intelligence and economic growth,” in *The economics of artificial intelligence: An agenda*, University of Chicago Press, 237–282.
- ALES, LAURENCE, CHRISTOPHE COMBEMALE, KATIE WHITEFOOT, AND ERICA FUCHS (2023): “How It’s Made: A General Theory of the Labor Implications of Technological Change,” Mimeo, Carnegie Mellon University.
- ATAK, JEREMY, ROBERT A. MARGO, AND PAUL W. RHODE (2019): ““Automation” of Manufacturing in the Late Nineteenth Century: The Hand and Machine Labor Study,” *Journal of Economic Perspectives*, 33, 51–70.
- AUTOR, DAVID, CAROLINE CHIN, ANNA M SALOMONS, AND BRYAN SEEGMILLER (2022): “New

- Frontiers: The Origins and Content of New Work, 1940–2018,” Working Paper 30389, National Bureau of Economic Research.
- AUTOR, DAVID, DAVID DORN, LAWRENCE F KATZ, CHRISTINA PATTERSON, AND JOHN VAN REENEN (2020): “The Fall of the Labor Share and the Rise of Superstar Firms,” *The Quarterly Journal of Economics*, 135, 645–709.
- AUTOR, DAVID H. (2019): “Work of the Past, Work of the Future,” *AEA Papers and Proceedings*, 109, 1–32.
- AUTOR, DAVID H, FRANK LEVY, AND RICHARD J MURNANE (2003): “The Skill Content of Recent Technological Change: An Empirical Exploration,” *The Quarterly Journal of Economics*, 118, 1279–1333.
- BESSEN, JAMES, MAARTEN GOOS, ANNA SALOMONS, AND WILJAN VAN DEN BERGE (2023): “What Happens to Workers at Firms that Automate?” *The Review of Economics and Statistics*, 1–45.
- BONFIGLIOLI, ALESSANDRA, ROSARIO CRINÒ, HARALD FADINGER, AND GINO GANCIA (2020): “Robot Imports and Firm-Level Outcomes,” Tech. rep., CEPR Discussion Papers No. 14593.
- BORUSYAK, KIRILL, RAFAEL DIX-CARNEIRO, AND BRIAN KOVAK (2022): “Understanding migration responses to local shocks,” Tech. rep., Mimeo, University of California, Berkeley.
- BOUSTAN, LEAH PLATT, JIWON CHOI, AND DAVID CLINGSMITH (2022): “Automation After the Assembly Line: Computerized Machine Tools, Employment and Productivity in the United States,” Working Paper 30400, National Bureau of Economic Research.
- BUERA, FRANCISCO J, JOSEPH P KABOSKI, RICHARD ROGERSON, AND JUAN I VIZCAINO (2021): “Skill-Biased Structural Change,” *The Review of Economic Studies*, 89, 592–625.
- CASELLI, FRANCESCO AND ALAN MANNING (2019): “Robot Arithmetic: New Technology and Wages,” *American Economic Review: Insights*, 1, 1–12.
- CHENG, HONG, LUKASZ A. DROZD, RAHUL GIRI, MATHIEU TASCHEREAU-DUMOUCHEL, AND JUNJIE XIA (2021): “The Future of Labor: Automation and the Labor Share in the Second Machine Age,” Tech. rep., FRB Philadelphia.
- COELLI, FEDERICA, DAVID DORN, DAVID HÉMOUS, AND MORTEN OLSEN (2023): “Automation and Employment,” Unpublished manuscript, University of Zurich, Department of Economics, Working Paper.
- CORTES, GUIDO MATIAS (2016): “Where Have the Middle-Wage Workers Gone? A Study of Polarization Using Panel Data,” *Journal of Labor Economics*, 34, 63–105.
- CURTIS, E MARK, DANIEL G GARRETT, ERIC C OHRN, KEVIN A ROBERTS, AND JUAN CARLOS SUÁREZ SERRATO (2021): “Capital investment and labor demand,” Tech. rep., National Bureau of Economic Research.
- DAO, MAI CHI, MITALI DAS, AND ZSOKA KOCZAN (2019): “Why is Labour Receiving a Smaller Share of Global Income?” *Economic Policy*, 34, 723–759.
- DAUTH, WOLFGANG, SEBASTIAN FINDEISEN, JENS SUEDEKUM, AND NICOLE WOESSNER (2021): “The adjustment of labor markets to robots,” *Journal of the European Economic Association*, 19, 3104–3153.

- DECHEZLEPRÊTRE, ANTOINE, DAVID HÉMOUS, MORTEN OLSEN, AND CARLO ZANELLA (2023): “Induced Automation Innovation: Evidence from Firm-level Patent Data,” *University of Zurich, Department of Economics, Working Paper*.
- DINLERSOZ, EMIN AND ZOLTAN WOLF (2023): “Automation, labor share, and productivity: Plant-level evidence from US Manufacturing,” *Economics of Innovation and New Technology*, 1–23.
- DONALD, ERIC (2022): “Optimal Taxation with Automation: Navigating Capital and Labor’s Complicated Relationship,” Mimeo, Boston University.
- ELOUNDOU, TYNA, SAM MANNING, PAMELA MISHKIN, AND DANIEL ROCK (2023): “Gpts are gpts: An early look at the labor market impact potential of large language models,” *arXiv preprint arXiv:2303.10130*.
- FEIGENBAUM, JAMES AND DANIEL P GROSS (2020): “Answering the Call of Automation: How the Labor Market Adjusted to the Mechanization of Telephone Operation,” Working Paper 28061, National Bureau of Economic Research.
- GOLDIN, CLAUDIA DALE AND LAWRENCE F KATZ (2008): *The Race Between Education and Technology*, Harvard University Press Cambridge.
- GOOS, MAARTEN, ALAN MANNING, AND ANNA SALOMONS (2014): “Explaining Job Polarization: Routine-Biased Technological Change and Offshoring,” *American Economic Review*, 104, 2509–26.
- GRAETZ, GEORG AND GUY MICHAELS (2018): “Robots at Work,” *The Review of Economics and Statistics*, 100, 753–768.
- GROSSMAN, GENE M AND EZRA OBERFIELD (2021): “The Elusive Explanation for the Declining Labor Share,” Working Paper 29165, National Bureau of Economic Research.
- GROSSMAN, GENE M. AND ESTEBAN ROSSI-HANSBERG (2008): “Trading Tasks: A Simple Theory of Offshoring,” *American Economic Review*, 98, 1978–97.
- HÉMOUS, DAVID AND MORTEN OLSEN (2022): “The Rise of the Machines: Automation, Horizontal Innovation, and Income Inequality,” *American Economic Journal: Macroeconomics*, 14, 179–223.
- HIRVONEN, JOHANNES, AAPO STENHAMMAR, AND JOONAS TUHKURI (2022): “New Evidence on the Effect of Technology on Employment and Skill Demand,” Tech. rep., Mimeo, Massachusetts Institute of Technology.
- HUBMER, JOACHIM (2023): “The Race Between Preferences and Technology,” *Econometrica*, 91, 227–261.
- HUBMER, JOACHIM AND PASCUAL RESTREPO (2021): “Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital-Labor Substitution,” Working Paper 28579, National Bureau of Economic Research.
- HUMLUM, ANDERS (2020): “Robot Adoption and Labor Market Dynamics,” Working paper, University of Chicago.
- JACKSON, MATTHEW AND ZAFER KANIK (2020): “How Automation that Substitutes for Labor Affects Production Networks, Growth, and Income Inequality,” Tech. rep., Stanford University.

- JONES, BENJAMIN F AND XIAOJIE LIU (2022): “A Framework for Economic Growth with Capital-Embodied Technical Change,” Working Paper 30459, National Bureau of Economic Research.
- KARABARBOUNIS, LOUKAS AND BRENT NEIMAN (2013): “The Global Decline of the Labor Share,” *The Quarterly Journal of Economics*, 129, 61–103.
- KEHRIG, MATTHIAS AND NICOLAS VINCENT (2021): “The Micro-Level Anatomy of the Labor Share Decline,” *The Quarterly Journal of Economics*, 136, 1031–1087.
- KIM, RYAN AND JONATHAN VOGEL (2021): “Trade Shocks and Labor Market Adjustment,” *American Economic Review: Insights*, 3, 115–30.
- KOCH, MICHAEL, ILYA MANUYLOV, AND MARCEL SMOLKA (2021): “Robots and Firms,” *The Economic Journal*, 131, 2553–2584.
- KOGAN, LEONID, DIMITRIS PAPANIKOLAOU, LAWRENCE D. W SCHMIDT, AND BRYAN SEEGMILLER (2021): “Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations,” Working Paper 29552, National Bureau of Economic Research.
- KRUSELL, PER, LEE E. OHANIAN, JOSÉ-VÍCTOR RÍOS-RULL, AND GIOVANNI L. VIOLANTE (2000): “Capital-Skill Complementarity and Inequality: A Macroeconomic Analysis,” *Econometrica*, 68, 1029–1053.
- MANN, KATJA AND LUKAS PUTTMANN (2023): “Benign Effects of Automation: New Evidence from Patent Texts,” *The Review of Economics and Statistics*, 105, 562–579.
- MARTINEZ, JOSEBA (2021): “Putty-Clay Automation,” Tech. rep., London Business School.
- MOLL, BENJAMIN, LUKASZ RACHEL, AND PASCUAL RESTREPO (2022): “Uneven Growth: Automation’s Impact on Income and Wealth Inequality,” *Econometrica*, 90, 2645–2683.
- NOY, SHAKKED AND WHITNEY ZHANG (2023): “Experimental evidence on the productivity effects of generative artificial intelligence,” *Science*, 381, 187–192.
- OBERFIELD, EZRA AND DEVESH RAVAL (2020): “Micro Data and Macro Technology,” *Econometrica*.
- OCAMPO, SERGIO (2022): “A task-based theory of occupations,” Tech. rep., Western University.
- PERETTO, PIETRO F. AND JOHN J. SEATER (2013): “Factor-eliminating technical change,” *Journal of Monetary Economics*, 60, 459–473.
- SIMON, HERBERT ALEXANDER (1965): *The shape of automation for men and management*, vol. 13, Harper & Row New York.
- TRAIBERMAN, SHARON (2019): “Occupations and Import Competition: Evidence from Denmark,” *American Economic Review*, 109, 4260–4301.
- WEBB, MICHAEL (2020): “The Impact of Artificial Intelligence on the Labor Market,” Working paper, Stanford University.
- ZEIRA, JOSEPH (1998): “Workers, Machines, and Economic Growth,” *The Quarterly Journal of Economics*, 113, 1091–1117.
- ZULETA, HERNANDO (2008): “Factor saving innovations and factor income shares,” *Review of Economic Dynamics*, 11, 836–851.

Supplementary Material: “Automation: Theory, Evidence, and Outlook”

Pascual Restrepo
December 23, 2023

A LABOR SUPPLY

This section microfounds the labor supply curve in the main text.

Assume there is a mass m_g of g households with utility

$$U_g = c_g - \frac{1}{1 + 1/\varepsilon} \cdot n_g^{1+1/\varepsilon},$$

where n_g are total hours worked and c_g is consumption per household. With these preferences, households' labor supply satisfies $w_g = n_g^{1/\varepsilon}$. This implies $\ell_g = m_g \cdot w_g^\varepsilon$.

B EQUILIBRIUM REPRESENTATION

This section shows that the equilibrium can be represented as the solution to the system of equations (1), (2), and (3). It also derives equation (4).

This result requires tasks to be assigned to a unique factor (except for zero-measure indifference sets). As in [Acemoglu and Restrepo \(2022\)](#), the following assumption ensures this is the case, and I impose this assumption in the rest of the appendix:

ASSUMPTION A1 (STRICT COMPARATIVE ADVANTAGE) *For all positive measure tasks sets $\mathcal{S} \subseteq \mathcal{T}$, $\psi_{gx}/(\psi_{kx} \cdot q_x)$ and ψ_{gx}/ψ_{jx} are not constant in \mathcal{S} for $j \neq g$.*

Equation (1) follows from equating the marginal cost of producing the final good, $c^f(p)$, to its price, which I normalized to 1.

For equation (2), let's consider the price index for product n . This is given by

$$p_n = \left(\int_{x \in \mathcal{T}_n} p_x^{1-\lambda} \cdot dx \right)^{\frac{1}{1-\lambda}}.$$

Tasks in $\mathcal{T}_{nk}(w)$ are produced with capital and have a price

$$p_x = 1/(\psi_{kx} \cdot q_x) \quad \text{if } x \in \mathcal{T}_{nk}(w).$$

Tasks in $\mathcal{T}_{ng}(w)$ are produced by g workers and have a price

$$p_x = w_g/\psi_{gx} \quad \text{if } x \in \mathcal{T}_{ng}(w).$$

Plugging these task prices in the price index for p_n yields (2).

For equation (3), let's consider the demand for labor coming from tasks x in $\mathcal{T}_{ng}(w)$, denoted by ℓ_{gx} . For each of these tasks, firms demand task x until

$$(A12) \quad p_x = p_n \cdot \left(\frac{y_n}{y_x} \right)^{-\frac{1}{\lambda}}.$$

For $x \in \mathcal{T}_{ng}(w)$, $p_x = w_g/\psi_{gx}$ and $y_x = \psi_{gx} \cdot \ell_{gx}$. Plugging in (A12) and solving for ℓ_{gx} yields:

$$(A13) \quad \ell_{gx} = y_n \cdot p_n^\lambda \cdot \psi_{gx}^{\lambda-1} \cdot w_g^{-\lambda}.$$

This can also be written as

$$\ell_{gx} = y \cdot s_y^n(p) \cdot p_n^{\lambda-1} \cdot \psi_{gx}^{\lambda-1} \cdot w_g^{-\lambda},$$

where $s_y^n(p) = p_n \cdot y_n/y = \partial \ln c^f(p)/\partial \ln p_n$ (from Shephard's lemma). Integrating this equation across tasks and products, labor market clearing becomes

$$y \cdot \sum_n s_y^n(p) \cdot p_n^{\lambda-1} \cdot \Gamma_{ng}(w) \cdot w_g^{-\lambda} = m_g \cdot w_g^\varepsilon.$$

Solving for w_g gives (3).

[Acemoglu and Restrepo \(2022\)](#) and [Acemoglu and Restrepo \(2023\)](#) provide conditions for the existence and uniqueness of a solution $\{y, w, p\}$ to this system of equations, as well as conditions under which the economy produces positive and finite output.

Finally, to derive Equation (4), integrate (A13) for tasks in $\mathcal{T}_{ng}(w)$. This yields

$$\ell_{ng} = y_n \cdot p_n^\lambda \cdot \Gamma_{ng}(w) \cdot w_g^{-\lambda} \Rightarrow s_n^g = \Gamma_{ng}(w)^{\frac{1}{\lambda}} \cdot \left(\frac{y_n}{\ell_{ng}} \right)^{\frac{1}{\lambda}-1},$$

where $s_n^g = w_g \cdot \ell_{ng}/(p_n \cdot y_n)$ denotes the share of group g wages in product n costs. The same steps but now applied to tasks in $\mathcal{T}_{nk}(w)$ imply

$$k_n = y_n \cdot p_n^\lambda \cdot \Gamma_{nk}(w) \Rightarrow s_n^k = \Gamma_{nk}(w)^{\frac{1}{\lambda}} \cdot \left(\frac{y_n}{k_n} \right)^{\frac{1}{\lambda}-1},$$

where $s_n^k = k_n/(p_n \cdot y_n)$ denotes the share of capital in product n costs.

Using the fact that $\sum_g s_n^g + s_n^k = 1$ (from constant returns to scale) yields

$$\sum_g \Gamma_{ng}(w)^{\frac{1}{\lambda}} \cdot \left(\frac{y_n}{\ell_{ng}}\right)^{\frac{1}{\lambda}-1} + \Gamma_{nk}(w)^{\frac{1}{\lambda}} \cdot \left(\frac{y_n}{k_n}\right)^{\frac{1}{\lambda}-1} = 1.$$

Using this equation to solve for y_n gives the formula in (4).

C AUTOMATION AND AN APPROXIMATION RESULT

This section provides a lemma that will be used extensively to derive the effects of automation. The lemma shows that, to a first-order approximation, one can decompose the impact of automation on task shares into the direct task displacement effects from automation and the endogenous reassignment of tasks.

LEMMA A1 (FIRST-ORDER EXPANSION OF TASK SHARES) *Consider an automation shock that automates tasks \mathcal{A}_{ng} across products and groups. Suppose that*

- i. the shock generates direct task displacements $\{d \ln \Gamma_{ng}^d\}$ across groups and products, with $d \ln \Gamma_{ng}^d < \epsilon$ for some $\epsilon > 0$.*
- ii. automated tasks \mathcal{A}_{ng} are in the interior of \mathcal{T}_{ng} ;*
- iii. the cost of producing these tasks with labor exceeds the cost of producing them with the newly developed capital, with the differences in cost exceeding δ for some $\delta > 0$.*

For small ϵ , the total effect of this shock on task shares can be approximated as

$$d \ln \Gamma_{ng} = -d \ln \Gamma_{ng}^d + \frac{\partial \ln \Gamma_{ng}(w)}{\partial \ln w} \cdot d \ln w + \mathcal{O}(\epsilon^2),$$

where $d \ln w$ denotes the equilibrium impact of the shock on wages.

PROOF. Let $w' = w + dw$ be the new equilibrium wages. Condition (i) implies that dw is $\mathcal{O}(\epsilon)$. Conditions (ii) and (iii) imply that, for small ϵ , all tasks in \mathcal{A}_{ng} will be automated at the new equilibrium wages w' and that these tasks will also be in the interior of $\mathcal{T}_{ng}(w')$.

These conditions then imply that the set of tasks performed by workers from group g

goes from $\mathcal{T}_{ng}(w)$ to $\mathcal{T}_{ng}(w') - \mathcal{A}_{ng}$, and the total change in Γ_{ng} is

$$d \ln \Gamma_{ng} = \frac{\int_{\mathcal{T}_{ng}(w')} \psi_{gx}^{\lambda-1} \cdot dx - \int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx} - \frac{\int_{\mathcal{A}_{ng}} \psi_{gx}^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx}.$$

By definition, the second term in this equation is $d \ln \Gamma_{ng}^d$. For the first term, perform a (log-linear) first-order Taylor expansion with respect to w (Assumption A1 implies task shares are differentiable functions of wages), which implies

$$\frac{\int_{\mathcal{T}_{ng}(w')} \psi_{gx}^{\lambda-1} \cdot dx - \int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx} = \frac{\partial \ln \Gamma_{ng}(w)}{d \ln w} \cdot d \ln w + \mathcal{O}(\epsilon^2),$$

establishing the lemma ■

D DERIVATIONS FOR THE GENERAL EQUILIBRIUM EFFECTS OF AUTOMATION

This section derives the GE effects of automation. These derivations provide first-order approximations to these effects valid when the fraction of automated tasks (measured by $d \ln \Gamma_{ng}^d$) are small and the conditions of Lemma A1 are met. This section derives expressions for the changes in product prices due to automation in equation (9), then turns to the dual version of Solow's residual in equation (8) and the expression for product n change in TFP, and concludes by deriving the equation for labor demand in equation (7).

Product n prices: Product n prices are given by (2). Using Lemma A1, express the GE effect of automation on product n prices as

$$d \ln p_n = \sum_g s_n^g \cdot d \ln w_g + \frac{1}{1-\lambda} \cdot \sum_g s_n^g \cdot \frac{\partial \ln \Gamma_{ng}}{\partial \ln w} \cdot d \ln w + \frac{1}{1-\lambda} \cdot \left[s_n^k \cdot d \ln \Gamma_{nk}^d - \sum_g s_n^g \cdot d \ln \Gamma_{ng}^d \right].$$

The first term captures the contribution of changes in input prices to costs, which depends on the share of wages in costs.

The second term captures the endogenous reallocation of tasks across skill groups in response to wage changes. This term adds up to zero since firms are indifferent between producing marginal tasks with g or j —an implication of the envelope theorem.

The third term captures the contribution of shifts in task shares due to automation. In

this term

$$d \ln \Gamma_{nk}^d = \sum_g \frac{\int_{\mathcal{A}_{ng}} [\psi_{kx} \cdot q'_x]^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{nk}(w)} [\psi_{kx} \cdot q'_x]^{\lambda-1} \cdot dx}$$

gives the expansion of capital's task share due to automation.

Rewrite $d \ln \Gamma_{nk}^d$ in terms of the sufficient statistics π_{ng} and $d \ln \Gamma_{ng}^d$ as

$$d \ln \Gamma_{nk}^d = \sum_g \frac{s_n^g}{s_n^k} \cdot d \ln \Gamma_{ng}^d \cdot [1 - (1 - \lambda) \cdot \pi_{ng}] \cdot dx,$$

with π_{ng} computed as an employment-weighted average of π_x over tasks in \mathcal{A}_{ng} .²⁷

$$\pi_{ng} = \int_{\mathcal{A}_{ng}} \frac{\psi_{gx}^{\lambda-1}}{\int_{\mathcal{A}_{ng}} \psi_g(\tilde{x})^{\lambda-1} \cdot d\tilde{x}} \cdot \pi(x) \cdot dx.$$

Using this expression for $d \ln \Gamma_{kn}$, express the change in prices as

$$(A14) \quad d \ln p_n = \sum_g s_n^g \cdot d \ln w_g - \sum_g s_n^g \cdot d \ln \Gamma_{ng}^d \cdot \pi_{ng},$$

which is equivalent to the expression in equation (9) in the text.

Changes in TFP: With constant returns to scale, changes in product n TFP satisfy²⁸

$$d \ln t f p_n = \sum_g s_n^g \cdot d \ln w_g - d \ln p_n.$$

²⁷The derivation follows from these algebraic manipulations:

$$\begin{aligned} d \ln \Gamma_{nk}^d &= \sum_g \frac{\Gamma_{ng}(w) \cdot w_g^{1-\lambda}}{\Gamma_{kn}(w)} \cdot \frac{\int_{\mathcal{A}_{ng}} \psi_{gx}^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_{ng}(w)} \psi_{gx}^{\lambda-1} \cdot dx} \cdot \frac{\int_{\mathcal{A}_{ng}} [\psi_{kx} \cdot q'_x]^{\lambda-1} \cdot dx}{\int_{\mathcal{A}_{ng}} \psi_{gx}^{\lambda-1} \cdot dx} \cdot w_g^{\lambda-1} \\ &= \sum_g \frac{s_n^g}{s_n^k} \cdot d \ln \Gamma_{ng}^d \cdot \frac{\int_{\mathcal{A}_{ng}} [\psi_{kx} \cdot q'_x]^{\lambda-1} \cdot dx}{\int_{\mathcal{A}_{ng}} \psi_{gx}^{\lambda-1} \cdot dx} \cdot w_g^{\lambda-1} \\ &= \sum_g \frac{s_n^g}{s_n^k} \cdot d \ln \Gamma_{ng}^d \cdot \int_{\mathcal{A}_{ng}} \frac{\psi_{gx}^{\lambda-1}}{\int_{\mathcal{A}_{ng}} \psi_{g\tilde{x}}^{\lambda-1} \cdot d\tilde{x}} \cdot \left[\frac{w_g \cdot \psi_{kx} \cdot q'_x}{\psi_{gx}} \right]^{\lambda-1} \cdot dx \\ &= \sum_g \frac{s_n^g}{s_n^k} \cdot d \ln \Gamma_{ng}^d \cdot \int_{\mathcal{A}_{ng}} \frac{\psi_{gx}^{\lambda-1}}{\int_{\mathcal{A}_{ng}} \psi_{g\tilde{x}}^{\lambda-1} \cdot d\tilde{x}} \cdot [1 - (1 - \lambda) \cdot \pi_x] \cdot dx \\ &= \sum_g \frac{s_n^g}{s_n^k} \cdot d \ln \Gamma_{ng}^d \cdot [1 - (1 - \lambda) \cdot \pi_{ng}] \cdot dx, \end{aligned}$$

²⁸To see this, write $p_n \cdot y_n = \sum_g \ell_{ng} \cdot w_g + k_n$. Totally differentiating both sides and rearranging yields $d \ln t f p_n = d \ln y_n - \sum_g s_n^g \cdot d \ln \ell_{ng} - s_n^k \cdot d \ln k_n = \sum_g s_n^g \cdot d \ln w_{ng} - d \ln p_n$.

Substituting for $d \ln p_n$ using (A14) gives the formula for product n TFP in the text.

To derive equation (8), differentiate the price-index condition in equation (1):

$$0 = \sum_n s_y^n \cdot d \ln p_n.$$

Substituting for $d \ln p_n$ using (A14) yields

$$\sum_n s_y^n \cdot \sum_g s_n^g \cdot d \ln w_g = \sum_n s_y^n \cdot \sum_g s_n^g \cdot d \ln \Gamma_{ng}^d \cdot \pi_{ng},$$

which is equivalent to (8), since $\sum_n s_y^n \cdot \sum_g s_n^g = s_y^g$.

Changes in wages and employment: Apply Lemma A1 to equation (3) to obtain

$$d \ln w_g = \frac{1}{\lambda + \varepsilon} \cdot d \ln y + \frac{1}{\lambda + \varepsilon} \cdot \sum_n \omega_g^n \cdot d \ln \zeta_n - \frac{1}{\lambda + \varepsilon} \cdot \sum_n \omega_g^n \cdot d \ln \Gamma_{ng}^d + \frac{1}{\lambda + \varepsilon} \cdot \mathcal{J}_\Gamma \cdot d \ln w.$$

This expression can be rewritten as

$$d \ln w_g = \frac{1}{\lambda + \varepsilon} \cdot d \ln y + \frac{1}{\lambda + \varepsilon} \cdot \sum_n (\omega_g^n - s_y^n) \cdot d \ln \zeta_n - \frac{1}{\lambda + \varepsilon} \cdot \sum_n \omega_g^n \cdot d \ln \Gamma_{ng}^d + \frac{1}{\lambda + \varepsilon} \cdot \mathcal{J}_\Gamma \cdot d \ln w.$$

by observing that

$$\sum_n s_y^n \cdot d \ln \zeta_n = (\lambda - 1) \cdot \sum_n s_y^n \cdot d \ln p_n + \sum_n s_y^n \cdot d \ln s_y^n = 0.$$

The solution to this fixed point problem gives (7). Using $d \ln \ell_g = \varepsilon \cdot d \ln w_g$ yields

$$(A15) \quad d \ln \ell_g = \varepsilon \cdot \Theta_g \cdot \text{stack}_j \left(\underbrace{d \ln y}_{\text{productivity effect}} + \underbrace{\sum_n (\omega_j^n - s_y^n) \cdot d \ln \zeta_n}_{\text{change in product mix}} - \underbrace{\sum_n \omega_j^n \cdot d \ln \Gamma_{nj}^d}_{\text{displacement effects}} \right).$$

GE effects on the aggregate labor share: The text describes a formula for the GE effects of automation on the aggregate labor share. This subsection derives this formula.

First, define $d \ln tfp = \sum_n s_y^n \cdot d \ln tfp_n$ as the aggregate TFP increase. Equation (8) implies that the percent increase in wages equals $d \ln w = \frac{d \ln tfp}{s_y^\ell}$, where s_y^ℓ is the aggregate labor share. The isoelastic labor supply implies a percent increase in labor income of

$$d \ln \left(\sum_g w_g \cdot \ell_g \right) = (1 + \varepsilon) \cdot \frac{d \ln tfp}{s_y^\ell}.$$

Let's compare this to the increase in output, $d \ln y$. Let $RS_g = \sum_j \theta_{gj}$ denote the row sums of the propagation matrix, and let $\bar{\Theta} = \sum_g s_y^g \cdot \Theta_g$. Solving for $d \ln y$ from (7) and (8):

$$d \ln y = \frac{1}{\sum_g s_y^g \cdot RS_g} \cdot \left[d \ln tfp - \bar{\Theta} \cdot \text{stack}_j \left(\sum_n (\omega_j^n - s_y^n) \cdot d \ln \zeta_n \right) + \bar{\Theta} \cdot \text{stack}_j \left(\sum_n \omega_j^n \cdot d \ln \Gamma_{nj}^d \right) \right].$$

Subtracting the change in output from the change in labor income, conclude that

$$\begin{aligned} d \ln s_y^\ell = & \left(1 + \varepsilon - \frac{s_y^\ell}{\sum_g s_y^g \cdot RS_g} \right) \cdot \frac{d \ln tfp}{s_y^\ell} \\ & + \frac{1}{\sum_g s_y^g \cdot RS_g} \cdot \left[\bar{\Theta} \cdot \text{stack}_j \left(\sum_n (\omega_j^n - s_y^n) \cdot d \ln \zeta_n \right) - \bar{\Theta} \cdot \text{stack}_j \left(\sum_n \omega_j^n \cdot d \ln \Gamma_{nj}^d \right) \right] \end{aligned}$$

The first line captures the GE effects on the aggregate labor share due to higher overall wages (as in [Grossman and Oberfield, 2021](#)).²⁹ The second line captures the contribution of reallocation across firms (as in [Oberfield and Raval, 2020](#)) or industries (as in [Acemoglu and Restrepo, 2022](#)) with different factor intensities. The last term is negative and captures the direct contribution of automation via extensive-margin changes in the task allocation.

E DERIVATIONS FOR THE DIRECT EFFECTS OF AUTOMATION

This section derives the direct effects of automation. I first explain the connection between my definition of direct effects and cross-sectional estimates.

Cross-sectional estimates and direct effects: In general, the demand for product n depends on y and the entire vector of product prices p . In my discussion, I assume

$$y_n = d_n(p_n, p_{idx}, y),$$

where p_{idx} is a common vector of moments of the price distribution. For example, if f is a CES, $y_n = y \cdot (p_n/p_{idx})^{-\varepsilon_d}$, where p_{idx} is the usual CES price index of all product prices. If f is a nested CES production function, one has a similar representation but now p_{idx} is the vector of price indices for all nests.

²⁹The term $1/RS_g - \varepsilon \geq \lambda$ provides a measure of how substitutable group g is for capital at the aggregate level. The (harmonic) mean $1/(\sum_g \omega_n^g \cdot RS_g) - \varepsilon \geq \lambda$ plays the same role as the aggregate elasticity of substitution in [Grossman and Oberfield \(2021\)](#): when it exceeds 1, shocks that increase mean real wages lower the labor share via this channel; when it is below 1, shocks that increase mean real wages raise the labor share via this channel.

Consider an outcome \mathcal{Y}_n for firm or product n . In equilibrium, this outcome is a function of output y , price statistics p_{idx} , factor prices w , and product n task shares $\Gamma_n(w)$. Write $\mathcal{Y}_n = G_n(y, p_{idx}, w, \Gamma_n(w))$ to indicate this dependency. In this representation, task shares summarize the technology for product n . Note that p_n is itself a function of wages and task shares $\Gamma_n(w)$, given in (2), and so it is not needed as an argument in G_n .

Consider the exogenous arrival of an automation shock that changes task shares by $d \ln \Gamma_n^d$, output by $d \ln y$, price statistics by $d \ln p_{idx}$, and wages by $d \ln w$. Using Lemma A1, write the impact of this shock on firm outcomes as

$$(A16) \quad d \ln \mathcal{Y}_n = \underbrace{\frac{\partial \ln G_n}{\partial \ln y} \cdot d \ln y + \frac{\partial \ln G_n}{\partial \ln p_{idx}} \cdot d \ln p_{idx} + \left[\frac{\partial \ln G_n}{\partial \ln w} + \frac{\partial \ln G_n}{\partial \ln \Gamma_n} \cdot \frac{\partial \ln \Gamma_n}{\partial \ln w} \right] \cdot d \ln w}_{\text{GE effects}} + \underbrace{\frac{\partial \ln G_n}{\partial \ln \Gamma_n} \cdot d \ln \Gamma_n^d}_{\text{direct effects, } \delta \ln \mathcal{Y}_n}.$$

The first line captures GE effects on product n due to expansions in the scale of the economy, changes in price indices affecting the demand for n , and changes in factor prices. The last term captures the direct effects of automation, defined as the change in product n outcomes holding y, p_{idx}, w constant. This gives the effect of automation working through changes in product n technology holding all aggregates and GE interactions fixed.

Well-identified empirical estimates compare units that satisfy the parallel trends assumption. This requires $\frac{\partial G_n}{\partial y}$, $\frac{\partial \ln G_n}{\partial \ln p_{idx}}$, and $\frac{\partial \ln G_n}{\partial \ln w} + \frac{\partial G_n}{\partial \Gamma_n} \cdot \frac{\partial \Gamma_n}{\partial w}$ to be equal across units. Otherwise, concurrent changes in wages, output, or the broader competitive environment would affect units differentially, violating parallel trends.

As a result, the GE effects in the first line of (A16) are absorbed by the constant term in these studies, and cross-sectional estimates identify the direct effects of automation. This observation motivates my definition of direct effects and my focus on these partial derivatives in the theory section.

Direct effects on product n prices: the GE effects on product n prices are given in (A14). From this equation, the direct effects of automation on prices are

$$(A17) \quad \delta \ln p_n = -s_n^\ell \cdot \sum_g \omega_n^g \cdot \pi_{ng} \cdot d \ln \Gamma_{ng}^d.$$

Note that this can also be written as $\delta \ln p_n = -d \ln t f p_n$.

Direct effects on product n labor share: the labor share in product n is

$$s_n^\ell = \frac{\sum_g w_g^{1-\lambda} \cdot \Gamma_{ng}(w)}{p_n^{1-\lambda}}.$$

The direct effects of automation are then given by

$$\delta \ln s_n^\ell = - \sum_g \omega_n^g \cdot d \ln \Gamma_{ng}^d - (1 - \lambda) \cdot \delta \ln p_n.$$

Replacing the formula for $\delta \ln p_n$ from equation (A17) gives equation (5) in the main text.

The text also describes the impact of automation on sales per worker, $p_n \cdot y_n / \ell_n$. This can be written as

$$\text{sales per worker}_n = \frac{1}{s_n^\ell} \cdot \bar{w}_n,$$

where \bar{w}_n denotes the average wage paid in product n . This decomposition implies

$$\delta \ln \text{sales per worker}_n = -\delta \ln s_n^\ell + \delta \ln \bar{w}_n,$$

which shows that automation can increase sales per worker via reductions in the labor share or by shifting the composition of the workforce towards higher-pay workers.

Direct effects on product n sales and employment: Recall that the demand for product n satisfies $y_n = d_n(p_n, p_{idx}, y)$. The direct effect of automation on the quantity of good n sold is therefore

$$\delta \ln y_n = -\varepsilon_n \cdot \delta \ln p_n = \varepsilon_n \cdot d \ln t f p_n,$$

where $\varepsilon_n = -\partial \ln d_n(p_n, p_{idx}, y) / \partial \ln p_n$ is the demand elasticity faced by firms producing n holding aggregate price indices and output constant. For example, if f were a CES production function, ε_n would equal the elasticity of substitution.

From this, the direct effect of automation on good n sales can be computed as

$$\delta(\ln p_n \cdot y_n) = (1 - \varepsilon_n) \cdot \delta \ln p_n = (\varepsilon_n - 1) \cdot d \ln t f p_n.$$

For employment, integrate equation (A13) for tasks in $\mathcal{T}_{ng}(w)$ to get

$$\ell_{ng} = y_n \cdot p_n^\lambda \cdot \Gamma_{ng}(w) \cdot w_g^{-\lambda}.$$

Partially differentiating both sides holding aggregates constant yields

$$\delta \ln \ell_{ng} = \delta \ln y_n + \lambda \cdot \delta \ln p_n - d \ln \Gamma_{ng}^d.$$

Substituting $\delta \ln y_n = \varepsilon_n \cdot d \ln t f p_n$ and $\delta \ln p_n = -d \ln t f p_n$ gives equation (11) in the text.

Direct effects on product n occupational structure: Denote by \mathcal{T}_{on} the set of tasks in product n that are part of occupation o , $\mathcal{T}_{ong}(w)$ the subset of these assigned to g , and \mathcal{A}_{ong} the subset of these that are automated. Finally, let $\Gamma_{ong}(w) = \int_{\mathcal{T}_{ong}(w)} \psi_{gx}^{\lambda-1} \cdot dx$ and $d \ln \Gamma_{ong}^d = \int_{\mathcal{A}_{ong}} \psi_{gx}^{\lambda-1} \cdot dx / \int_{\mathcal{T}_{ong}(w)} \psi_{gx}^{\lambda-1} \cdot dx$.

Using this notation, wage payments to group g in product n in occupation o can be computed by integrating equation (A13) for tasks in $\mathcal{T}_{ong}(w)$:

$$w_g \cdot \ell_{ong} = y_n \cdot p_n^\lambda \cdot \Gamma_{ong}(w) \cdot w_g^{1-\lambda}.$$

Total wage payments for occupation o in product n are then equal to

$$\sum_g w_g \cdot \ell_{ong} = y_n \cdot p_n^\lambda \cdot \sum_g \Gamma_{ong}(w) \cdot w_g^{1-\lambda}.$$

Partially differentiating both sides holding aggregates constant yields

$$\delta \ln \left(\sum_g w_g \cdot \ell_{ong} \right) = - \sum_g \omega_{on}^g d \ln \Gamma_{ong}^d.$$

Equation (6) follows from $\delta \ln \frac{\omega_n^o}{\omega_n^g} = \delta \ln \left(\sum_g w_g \cdot \ell_{ong} \right) - \delta \ln \left(\sum_g w_g \cdot \ell_{o'ng} \right)$.

F THE ELASTICITY OF SUBSTITUTION BETWEEN CAPITAL AND LABOR

The text provides a formula for the elasticity of substitution between capital and labor in product n . This elasticity is defined as the percent increase in capital-to-labor inputs

following a common increase in wages of $d \ln w_g = d \ln w$, or

$$\sigma_n = \frac{d \ln k_n / d \ln \ell_n}{d \ln w}.$$

This corresponds to the elasticity one would identify by exploiting differences in wages faced by firms, as in (Oberfield and Raval, 2020). This definition can be rewritten in terms of the response of the labor share in product n in response to a common increase in wages:

$$\sigma_n = 1 - \frac{1}{1 - s_n^\ell} \cdot \frac{d \ln s_n^\ell}{d \ln w}.$$

To compute this elasticity, use the fact that the labor share in product n is

$$s_n^\ell = p_n^{\lambda-1} \cdot \sum_g \Gamma_{ng}(w) \cdot w_g^{1-\lambda}.$$

The change in s_n^ℓ from a common increase in wages $d \ln w$ is

$$\frac{d \ln s_n^\ell}{d \ln w} = (\lambda - 1) \cdot s_n^\ell + (1 - \lambda) + \sum_g \omega_n^g \cdot \sum_j \frac{\partial \ln \Gamma_{ng}(w)}{\partial \ln w_j}.$$

Plugging this expression in the definition of σ_n yields the formula in the text.