

Chapter 1

Tasks at work: comparative advantage, technology and labor demand[☆]

Daron Acemoglu^{a,*}, Fredric Kong^a, and Pascual Restrepo^b

^aMIT, Cambridge, MA, United States, ^bYale University, New Haven, CT, United States

*Corresponding author. e-mail address: daron@mit.edu

Contents

1	Introduction	2	5	Putting it all together: shocks and propagation in the one-sector economy	32	
	Chapter outline	5		5.1 Equilibrium: ripple effects and the propagation matrix	33	
	Tasks: a partial review of the theoretical literature	7		5.2 Properties of the propagation matrix	33	
	Tasks: a partial review of the empirical literature	9		5.3 Automation	36	
				5.4 New tasks	39	
				5.5 Labor-augmenting technology	39	
				5.6 Capital-augmenting technologies	41	
				5.7 Changes in labor supply	42	
2	The task model: the one-sector case	10		6	The multi-sector economy	43
	2.1 Environment	11		6.1 Environment	43	
	2.2 Equilibrium	13		6.2 Equilibrium	44	
	2.3 Equilibrium representation in terms of task shares	14		6.3 Effects of technology in the multi-sector economy	46	
	2.4 Beyond CES	16		6.4 Sectoral TFP and markups	49	
3	Different technology, different effects	17		7	Reduced-form evidence	50
	3.1 The no-ripples economy	17		7.1 US labor market trends	50	
	3.2 Automation	18		7.2 Specification	52	
	3.3 New tasks	22		7.3 Measuring automation and new tasks	55	
	3.4 Labor-augmenting technologies	24		7.4 Reduced-form estimates	59	
	3.5 Capital-augmenting technologies	27		7.5 Robustness	64	
	3.6 Microfoundation for shifting Cobb-Douglas exponents	27		7.6 Taking stock	66	
	3.7 Taking stock	28				
4	From micro to macro elasticities	28				
	4.1 Macroeconomic elasticities of substitution	28				
	4.2 Examples	30				

[☆] We thank the editors Christian Dustmann and Thomas Lemieux for their comments. Acemoglu gratefully acknowledges support from the Hewlett and Smith Richardson Foundations, Kong from the Bradley Foundation, and Restrepo from the National Science Foundation (Award No. 2049427).

8 Estimation of general equilibrium effects and counterfactuals	66	B.2 Effects of technology with ripples	98
8.1 General equilibrium effects of technology and markups	67	Appendix C Equilibrium in the multi-sector economy	99
8.2 Estimating the propagation matrix	69	Appendix D Endogenous labor supply	101
8.3 Decompositions	74	Appendix E Derivations for the Allen-Uzawa elasticities of substitution and properties of the propagation matrix	103
8.4 Limited distributional impacts of labor-augmenting technologies	79		
8.5 The missing technology puzzle	80	Appendix F Additional empirical results	106
8.6 The incidence puzzle	82	F.1 Robustness checks	106
9 Conclusion	83	F.2 Estimating the propagation matrix	107
Appendix A Equilibrium existence and uniqueness	87	References	111
Appendix B Effects of technology	89		
B.1 No-ripple economy	94		

1 Introduction

The wage and occupational structures of the United States and other industrialized countries have experienced epochal changes over the past several decades. US wage inequality has soared, while the real wages of less-educated workers have stagnated or fallen, and their employment rates have declined. Simultaneously, employment has shifted from production and clerical occupations to higher-paying managerial, professional, and technical jobs and various service occupations with lower pay. These trends have been accompanied by a lower labor share, especially in manufacturing, and lackluster productivity growth.¹ Early research explored the contribution of labor demand to these trends using a (reduced-form) approach based on an aggregate production function and technologies that augment skilled or unskilled labor.² In this canonical approach, labor demand changes are combined with labor supply and institutional factors to account for the observed trends.

A more recent strand departs from this approach and starts with a setup in which the production of goods and services requires the completion of tasks, and factors of production are assigned to perform these tasks.³ For example,

¹ For a summary of the wage and inequality trends, see Goldin and Katz (2008), Acemoglu and Autor (2011), Acemoglu and Restrepo (2019), Autor (2019), Restrepo (2024). Karabarbounis and Neiman (2013) documents the decline in the labor share in the United States and other industrialized countries, while Acemoglu and Autor (2011) and Goos et al. (2014) show correlated shifts in occupational structure across several OECD economies. For recent reviews of trends in the wage structure in European and OECD countries, see, e.g., Gornick (2024).

² See, among others, Bound and Johnson (1992), Katz and Murphy (1992), Berman et al. (1994) and Autor et al. (1998). See Acemoglu (2002) for a review and extensions of these approaches.

³ See Autor et al. (2003), Acemoglu and Autor (2011) and Autor and Handel (2013) for some of the early works using the task approach to study inequality. We discuss the evolution of this literature at the end of the Introduction.

the production of a smartphone relies on a range of design and planning tasks; the manufacturing of the microchip, the battery, the camera, the speakers, the screen, numerous different types of sensors, and various other components; assembly of these components; and a series of non-production tasks, including various back-office activities, quality control, and inventory management. In addition, several marketing, advertising, transport, wholesale, and retail functions must be completed for smartphones to reach consumers. Each task needs to be assigned to various factors of production. For example, assembly can be performed by craft workers, low-skill workers, a combination of computerized equipment and human labor, or by robots.

In this *task framework*, the assignment of tasks to factors is shaped by technology and mediates the effect of technology on productivity and wages. For example, the task assignment depends on whether some tasks are standardized and can be performed by unskilled labor or whether technology permits the tasks to be performed by machines or algorithms. Technological change can then significantly impact productivity and equilibrium factor prices by enabling new ways of completing tasks. This can happen via *automation*, which occurs when new equipment, robots, software, or algorithms take over tasks previously performed by labor, as well as via *new tasks*, which entails the introduction of new tasks performed by labor.

The task framework is useful not only because it brings greater descriptive realism to modeling the production process but also because it generates a more comprehensive set of comparative statics regarding the impact of different technological advances and allows for richer substitution patterns between factors of production that shape their (general) equilibrium effects.

Different technologies, different effects The early literature on wage inequality in labor and macroeconomics assumed that all technologies work by augmenting factors of production, increasing the quantity or quality of their output. This restrictive view of technology drove some of its major conclusions. For example, an implication of the standard models discussed in Acemoglu (2002) is that skill-biased technological change (modeled as an increase in the productivity of skilled workers) always raises the real wages of low-skill workers, even as it increases inequality.⁴

In reality, technologies take more variegated forms and have richer effects on wages, inequality, and productivity. Besides augmenting workers or capital uniformly at all tasks, new technologies can:

- Increase workers' productivity in some tasks currently assigned to them. For example, a better drill makes workers more productive at drilling but not at other tasks. This type of *labor-augmenting* change occurs at the intensive margin. Our framework shows that this form of technology generates relatively small effects on wages and inequality and ambiguous impacts on the labor share of national income.

⁴ See Acemoglu and Autor (2011) for other implications that follow from earlier modeling assumptions.

- Increase capital productivity in some tasks currently assigned to capital. An example is a new and more powerful software system replacing older inventory management methods. This type of *capital-augmenting* change at the intensive margin raises productivity and always pushes up real wages but has ambiguous and minor effects on the labor share.

More novel and unique to the task framework, new technologies can also:

- *Automate work.* New technologies achieve this by enabling the use of equipment, software, and algorithms to perform tasks previously assigned to labor. Examples include software systems that take over office tasks previously assigned to workers or robots that now perform welding, cutting, painting, and assembly tasks. Automation can have major distributional effects, while its productivity impacts can be limited. Moreover, automation always reduces the labor share and can depress the real wages of displaced workers.
- *Create new tasks.* New tasks increase productivity by reorganizing production or introducing a finer division of labor. New tasks assigned to labor tend to raise the real wages of all skill groups and the labor share of national income. Computer-assisted design tools, machinery that enable novel technical work, and new programming, integration, and customer service functions introduced by recent technologies are examples of new tasks.

The discussion above showcases a key insight from the task framework: different technologies have different impacts. For example, labor-augmenting technology and new tasks can have opposite effects. Technologies that augment labor in some of their current tasks can reduce the real wages of affected groups, especially if the demand for these tasks is inelastic. In contrast, technologies that create new tasks for workers always increase their wages and raise the labor share. This critical distinction argues against the use of “augmenting technology” as a catchphrase for all technologies that work with labor. It also argues against the presumption that a technology that “augments” workers in some of their tasks necessarily raises their wages.

Flexible substitution between factors depending on comparative advantage
 Our framework distinguishes between microeconomic and macroeconomic substitution. Even though different workers and capital are perfect substitutes in producing a given task (at the micro level), they are imperfect substitutes at the aggregate level because they specialize in different tasks according to their comparative advantage. The aggregate substitution patterns depend on the strength of comparative advantage and the extent to which groups compete for marginal tasks, generating rich aggregate substitution patterns between factors.

These aggregate substitution patterns are essential for understanding the equilibrium effects of technology. Consider, for example, the automation of tasks performed by a group of workers. This shock reduces the group’s relative wage, creating an endogenous reassignment of marginal tasks toward this group. This affects other workers’ wages and creates further reassessments. The

strength of these *ripple effects* depends on the aggregate substitution patterns between groups. We show that ripple effects can be summarized (up to a first order) by a *propagation matrix*. This matrix determines how shifts in demand or supply impact wages, not only for workers directly exposed to the change in technology but also for workers competing against them for tasks. The propagation matrix captures the intuitive idea that a shock affecting one group generates an indirect impact on other highly substitutable groups.

Besides these conceptual innovations, the task model provides tractable equations that describe how changes in group-level wages depend on advances in different types of technologies and other labor demand forces, such as offshoring, structural change, and product market structure, including markups. These equations can be further extended to account for institutional and supply-side factors.

The equations decompose the impact of demand-side forces into a productivity effect; measures of the direct effects of technology on labor demand (e.g., the reallocation of tasks from labor to capital because of automation or the increased demand for labor in new tasks); a term capturing shifts in the economy's sectoral composition; and ripple effects summarized by the propagation matrix. This decomposition illustrates the channels through which technology affects wages. For example, automation impacts labor demand mainly by reallocating tasks from labor to capital. Instead, industry-level productivity shocks influence labor demand mainly by shifting the economy's sectoral composition.

Moreover, this characterization can be used to derive simple reduced-form equations or to conduct structural exercises quantifying the contribution of different demand forces to observed changes in the wage structure. We demonstrate both uses with an application to US data.

Chapter outline

This chapter reviews recent advances in the task framework and shows how this framework can be a powerful tool for theoretical, reduced-form, and structural research. The first part of the chapter introduces the task framework, explains its distinguishing features, derives the equations for wage changes, and presents a range of comparative statics describing the effects of technology on wages and factor shares. This part of the chapter builds on Acemoglu and Restrepo (2022). The new element is drawing out the implications of new tasks for the wage and employment structure of the economy, which has not been the focus of past work.

Section 2 introduces a one-sector version of the task framework with multiple types of skills, tasks, and technologies, and defines and characterizes the competitive equilibrium in this economy.

Section 3 specializes this environment to what we call the “no-ripples economy” to provide a transparent exposition of the varying effects of different types

of technologies. This example economy shuts down the endogenous reassignment of tasks across worker groups and the resulting ripple effects.

Section 4 clarifies the distinction between microeconomic and macroeconomic elasticities of substitution and how the latter elasticity is shaped by competition for marginal tasks and comparative advantage schedules.

Section 5 introduces the propagation matrix, which summarizes the rich substitution patterns implied by the task framework and uses this matrix to provide a full characterization of the equilibrium, including the ripple effects.

Section 6 extends this economy to a multi-sector economy, which is the basis of our empirical application. This section also introduces product market markups and characterizes their impact on the wage structure.

In the second part of the chapter, we use the wage equations derived from the task model to conduct a reduced-form analysis and then a structural exercise quantifying the importance of automation, new tasks, and other forces to the observed changes in the wage structure.

Section 7 derives simple reduced-form equations that relate wage changes across different worker groups to measures of the direct impact of automation, new tasks, markups, sectoral TFP changes, and labor-augmenting technologies. We estimate these reduced-form equations using publicly available US data. In particular, we use data from 500 groups of US workers, defined by education, gender, age, race, and native/foreign-born status, as our skill groups and focus on changes from 1980 to 2016. As part of this exercise, we introduce a new measure of new tasks across these groups. This part of the chapter also draws on past work, but the estimation of the effects of new tasks are original to this chapter.

We document that a 10% loss of tasks for a group due to automation during this period leads to a 12% relative wage decline and 8.2% reduction in hours worked per person. Using the measure of new tasks, we document that 10% additional new tasks for a group lead to an 8.5% increase in relative wage and 26% increase in hours worked per person. Overall, in the reduced form, the change in the share of tasks across groups due to automation and new tasks accounts for 67%-84% of the changes in the between-group wage structure in the US during this period and 53%-68% of the changes in group-level employment. We also estimate the reduced-form distributional effects of other factors, including sectoral reallocation, sectoral TFP trends, and changes in product market markups. These factors appear to have played a more limited role in the changes in the US wage structure. For example, while automation and new tasks jointly explain 67%-84% of the variation in between-group wage growth in the US from 1980 to 2016, proxies for skill-biased factor-augmenting technologies explain no more than a few percentage points of these changes.

The entire real wage impacts of technology cannot be estimated using these reduced-form equations because the constants in the reduced-form equations absorb their productivity effects and because potentially complex ripple effects are ignored. In Section 8, we outline a tractable structural approach for estimating

the first-order effects of automation, new tasks, and other shocks and estimate these equilibrium effects. The approach uses the equations for equilibrium wage changes described above combined with measures of automation and new tasks, our estimates of the propagation matrix, and existing estimates of the elasticities of substitution between industries and between tasks. This method allows us to quantify the full general equilibrium impacts of automation and new tasks and conduct counterfactual analyses.

Section 9 concludes and proposes areas for future research. The Appendix contains proofs, theoretical derivations, and additional empirical results.

Tasks: a partial review of the theoretical literature

The microfoundations of the task model go back to Zeira (1998), who considers a model where aggregate output is produced from a continuum of product lines (similar to tasks here), which can be allocated to capital or labor. Economic growth is driven by innovations that reallocate product lines/tasks away from labor toward capital.

Acemoglu and Zilibotti (2001) build a model in which two types of labor have different comparative advantages across a continuum of tasks, and technology affects the task production functions. This model is used to study how new technologies developed in the industrialized world influence inequality and growth in these economies as well as in developing countries, and especially how these technologies may be inappropriate for the needs of developing economies.

The first paper to use the task framework for systematically analyzing inequality is Autor et al. (2003). This paper builds a model with three tasks—one that corresponds to nonroutine problem-solving and complex communication activities performed by skilled labor, one that corresponds to nonroutine manual work performed by unskilled labor, and one that is closely associated with routine cognitive and manual tasks. The authors argue that computers can replace workers engaged in routine cognitive and manual activities because they can cheaply perform routine tasks that can be codified into step-by-step instructions. Computers can also, directly and indirectly, complement workers in nonroutine problem-solving and complex communications tasks. These authors develop a novel empirical mapping from these tasks to data and undertake the first comprehensive empirical analysis of the implications of the task model. Autor and Handel (2013) further extend both the theoretical framework and the measurement of the task content of occupations.

Acemoglu and Autor (2011) build a model that combines elements from the papers mentioned above and the classic Ricardian trade framework of Dornbusch et al. (1977). In their model, there are three types of workers (low, middle and high skill) and a continuum of tasks. Higher-skilled workers are assumed to have a comparative advantage in higher-indexed (more complex) tasks. Technological change can augment one or multiple labor types, and enables the

automation of some tasks using new equipment or software. This paper clarifies the distinction between standard (factor-augmenting) skill-biased technological change and automation—emphasizing how these technologies impact different parts of the earnings distribution and can have distinct effects on the level of real wages and inequality. This work also highlights the connection between the task framework and the earlier assignment literature for example, how the task approach builds on the competitive assignment setup of Sattinger (1975) and Teulings (1995, 2005) as well as the international trade literature focusing on offshoring of tasks, such as Grossman and Rossi-Hansberg (2008), Rodríguez-Clare (2010) and Acemoglu et al. (2015).

Our approach in this chapter builds more directly on recent work in task-based models. Acemoglu and Restrepo (2018b) develop a tractable task-based model and generalize this framework by introducing new tasks. This paper also demonstrates how the combination of automation and new tasks can lead to balanced economic growth, provided that the decline in the labor share and the contraction in the range of tasks induced by automation need to be compensated by creating new (labor-intensive) tasks. Acemoglu and Restrepo (2020b) extend this framework and draw the implications of automation and new tasks for wage inequality.

Acemoglu and Restrepo (2020a) use a task model to study the implications of industrial robot adoption in US manufacturing. Their work shows how simple estimating equations can be derived from the task model. Their estimates show that industrial robots impacted wages and employment, especially for workers specializing in manual blue-collar tasks in local labor markets exposed to these new technologies. This work also clarifies how the aggregate effects of this type of automation can be computed by combining the productivity impacts of robots with reduced-form estimates of the displacement effects.

Our treatment in this chapter builds most closely on Acemoglu and Restrepo (2022). This paper introduces a general version of the task model with multiple skill groups and with a flexible pattern of comparative advantage. Despite the generality of the model, the paper shows that the equilibrium takes a simple form and enables the empirical exploration of the consequences of different technologies and their propagation. This paper further clarifies the distinction between capital-skill complementarity, which increases the quantity or quality of capital as discussed by Griliches (1969), Berman et al. (1994), and Krusell et al. (2000), and automation, which is driven by improvements in capital productivity for tasks previously performed by labor. While the former process affects inequality indirectly—by increasing the output of capital-intensive activities or sectors—automation impacts inequality directly by displacing some groups of workers from the tasks they used to perform.

Other contributions exploring the implications of automation in task-based models include Acemoglu and Restrepo (2018a), Acemoglu and Restrepo (2019), Aghion et al. (2018), Feng and Graetz (2020), Moll et al. (2022), Nakamura and Zeira (2024), Jones and Liu (2022), Hubmer and Restrepo (2021) and

Acemoglu and Loebbing (2024). Another branch of the literature proposes models of factor-eliminating technical change, where technology works by reducing the weight of a factor in the production process (see, for example, Zuleta, 2008; Peretto and Seater, 2013). We show below that the task framework provides a microfoundation for this form of technological progress.

Tasks: a partial review of the empirical literature

An active and growing empirical literature has explored the implications of automation and new tasks for the wage structure. This literature is surveyed in Restrepo (2024). Much of this literature focuses on the US and finds evidence that automation technologies reduce the labor share (and increase sales per worker), for example, see Acemoglu and Restrepo (2020a) for the effects of industrial robots across industries and local labor markets, and Boustan et al. (2022) for the effects of CNC technologies in US manufacturing. Kogan et al. (2021), Dechezleprêtre et al. (2023) and Autor et al. (2024) report a negative association between the deployment of automation technologies measured using patent data and the labor share across US industries and occupations, while an extensive literature building on Autor et al. (2003) documents a negative relationship between automation and employment in routine jobs (see, for example, Webb, 2020; Kogan et al., 2021). Autor et al. (2024) additionally show that occupations experiencing the introduction of new tasks expanded their employment.

We see similar patterns beyond the US. Several industrial economies have experienced declining labor shares since the 1980s, especially in manufacturing (Karabarbounis and Neiman, 2013) and a declining share of employment in routine occupations (Goos and Manning, 2007; Acemoglu and Autor, 2011; Goos et al., 2014)—both telltale signs of automation. Consistent with this interpretation, Graetz and Michaels (2018) document a link between robot adoption and labor share changes by exploiting cross-country and cross-industry variation. A growing literature using firm-level data on robot adoption across a wide range of countries, including Denmark (Humlum, 2020), France (Bonfiglioli et al., 2020; Acemoglu et al., 2020), and the Netherlands (Acemoglu et al., 2023) finds that robot adoption is associated with a reduction in labor shares and the share of employment in routine jobs, in line with the predictions of the task model. Acemoglu et al. (2023) also show that workers specialized in blue-collar, routine tasks are the ones that are negatively impacted by robots, as predicted by the task framework.

Concurrently, we see rising wage inequality in several, though not all, industrialized countries. The college premium rose in the US, Canada, Mexico, Japan, the UK and Sweden; remained stable in France, Italy and Russia; and actually decreased in Korea, Netherlands and Spain (see Katz and Autor, 1999; Krueger et al., 2010). The increase in wage inequality is more pervasive when focusing on the difference in wages between the 90th and 10th percentile or the

total variance of log wages. For example, total variance of log wages increased in the US, UK, Canada, Germany, Italy and Mexico, but decreased in Russia, Spain and Sweden (see Krueger et al., 2010). Similarly, Machin and Van Reenen (2010) and Van Reenen (2011) document growing 90-10 male wage inequality in Denmark, Japan, Netherlands, New Zealand, the UK and the US from 1980 to 1990. Since 1990, we have also seen rising 90-10 inequality in Australia, Finland, Germany and Sweden (France being the only country in their sample where 90-10 inequality appears not to have increased).

The German case is particularly interesting. The comprehensive study by Dustmann et al. (2009) documents an increase in wage inequality in West Germany (measured by the dispersion in log wages) dating back to the 1970s for men and to the 1990s for women. The authors also show that the 85th percentile of wages for both men and women rose more rapidly than median wages or wages at the 15th percentile from 1975 to 2004. Wages at the bottom have stagnated or decreased since the early 1990s. Simultaneously, the premium paid to workers with an apprenticeship or college degree relative to those with no post-secondary schooling rose, while the premium earned by college graduates relative to workers with apprenticeship has remained stable.⁵

Overall, even though there is evidence of higher wage inequality in some European economies, the increase has been less pronounced and pervasive than in the US. One possibility is that these divergent experiences are due to differences in European labor market institutions that generate wage compression and limit the response of wages to changes in technology. For example, Cahuc (2024) argues that a high minimum wage and rigid wage structure have kept inequality in check in France, but this came at the expense of growing disparities in employment rates between more and less educated workers. In light of the existing evidence, it is therefore reasonable to conjecture that automation could have been a source of declining labor shares and rising inequalities in other industrialized economies as well, but we are not aware of systematic analyses of the effects of automation (or new tasks) on inequality in Europe. Any such study may have to incorporate the influence of different labor market institutions on wage and employment responses.

2 The task model: the one-sector case

This section introduces the task model and characterizes the equilibrium. We focus on the one-sector version of the model for simplicity, returning to the multi-sector economy in Section 6.

⁵ Dustmann et al. (2009) also perform an accounting exercise that removes the influence of changes in the supply of skills using the methodology of Katz and Murphy (1992). They find evidence of a rising relative demand for education, though these changes are less pronounced than those for the US. It is noteworthy that Dustmann et al. (2009) use the IABS dataset. Other studies using the GSOEP, including Fuchs-Schündeln et al. (2010), find a modest increase in total log wage variance and no evidence of a rising college premium (though their analysis also pools apprentices and workers with no post-secondary education together, rather than separating them as in Dustmann et al. (2009)).

2.1 Environment

A (unique) final good y is produced by combining a set of complementary tasks $x \in \mathcal{T}$ with measure $M > 0$. This good is set as the numeraire, with price normalized to 1. Task quantities $y(x)$ are aggregated using a constant elasticity of substitution (CES) aggregator with elasticity $\lambda \in (0, 1)$,

$$y = \left(\frac{1}{M} \int_{\mathcal{T}} (M \cdot y(x))^{\frac{\lambda-1}{\lambda}} dx \right)^{\frac{\lambda}{\lambda-1}}.$$

The set \mathcal{T} is assumed measurable and “ dx ” denotes the Lebesgue integral. \mathcal{T} could represent a continuum of tasks arranged along a line (as in Acemoglu and Autor, 2011), or could be a region of the plane or a multi-dimensional space.

The key economic decision in this model is the allocation of the tasks in \mathcal{T} to factors of production. The total quantity produced of task x is assumed to be

$$y(x) = A_k \cdot \psi_k(x) \cdot k(x) + \sum_g A_g \cdot \psi_g(x) \cdot \ell_g(x). \quad (1)$$

Intuitively, tasks can be produced by workers of different skill types, indexed by $g \in \mathbb{G} = \{1, 2, \dots, G\}$ or by (specialized) capital equipment. We denote the quantity of labor of skill type g used in task x by $\ell_g(x)$ and the amount of capital used in the production of task x by $k(x)$. Workers in skill group g have productivity $A_g \cdot \psi_g(x) \geq 0$ in task x , where the $\psi_g(x)$ schedule represents their comparative advantage across tasks. Capital has productivity $A_k \cdot \psi_k(x) \geq 0$ in task x , which is equal to zero for tasks where technology does not yet permit capital to substitute for workers. The A_k and A_g terms represent standard factor-augmenting technologies, which make factors uniformly more productive in all tasks.

Eq. (1) imposes perfect substitutability of capital and the different groups of workers at the task level. This feature of the model is a simplifying, but not implausible, assumption. Many new equipment and software types, such as computer numerical control machinery and robots, can perform various tasks with little human involvement (while the programming, maintenance, and service of such equipment correspond to other tasks that remain labor-intensive). This feature is a simplification since some labor-intensive tasks require tools (e.g., hammers), but it does not affect the implications of the framework.⁶

Labor supply is assumed inelastic, with the total supply of group g denoted as ℓ_g , while the real wage of this group is denoted by w_g . We discuss elastic labor supply in Section 8.

To keep the model static, capital is treated as an intermediate good, produced using units of the final good and used up in the same period due to depreciation.

⁶ It is straightforward to generalize this production function so that labor uses some tools and capital equipment needs operators. So long as the share of these factors is small, all implications of our framework continue to hold. See the discussion in the online appendix of Acemoglu and Restrepo (2018b).

Specifically, capital of type x , $k(x)$, is produced using the final good at a constant marginal cost normalized to 1. Changes in the productivity and cost of capital are subsumed into changes in the $\psi_k(x)$ schedules. Net output, which is equal to consumption, is therefore obtained by subtracting the production cost of capital goods from output:

$$c = y - \int_{\mathcal{T}} k(x) \cdot dx.$$

Following Acemoglu and Restrepo (2022), throughout we impose the following restrictions on the task space, which are sufficient for the existence of a unique equilibrium where all workers are assigned a positive measure of tasks and output is positive and finite. While these assumptions can be weakened, this would be at the cost of additional complications and we do not pursue this path here.

Assumption 1 (Restrictions on the task space). • For each task $x \in \mathcal{T}$, there exists at least one $g \in \mathbb{G}$ such that $\psi_g(x) > 0$. Moreover, the integrals

$$\int_{x: \psi_g(x) > 0} \psi_g(x)^{\lambda-1} \cdot dx$$

are finite.

- For each $g \in \mathbb{G}$, there is a positive measure of tasks x for which $\psi_g(x) > 0$, $\psi_{g'}(x) = 0$ for all other $g' \neq g$, and $\psi_k(x) = 0$.
- Comparative advantage is strict. For any two groups $g \neq g'$ and constant $a > 0$, the set of tasks such that $\psi_g(x)/\psi_{g'}(x) = a$ has measure zero. For any group g and constant $a > 0$, the set of tasks such that $\psi_g(x)/\psi_k(x) = a$ has measure zero.

Part 1 of the assumption is a sufficient condition for positive output in the economy (otherwise, such an economy may generate zero output). Part 2 guarantees that all skill groups are necessary for production and implies that technological changes will not make any skill group completely redundant. These conditions also ensure that output is always finite (because it rules out the possibility that capital will perform all tasks). Part 3 of the assumption imposes strict comparative advantage. This removes any indeterminacy in the allocation of tasks to workers and ensures that ties (situations in which a task can be produced in a cost-minimizing way with more than one factor) occur only on measure zero sets. Throughout, we also adopt the (non-consequential) tie-breaking rule that whenever there is a tie, tasks are allocated to capital first and then to lower-indexed skill types ahead of higher-indexed skill types.

2.2 Equilibrium

A market equilibrium is defined by a positive vector of real wages $w = \{w_g\}_{g \in \mathbb{G}}$, an output level y , an allocation of tasks to worker groups $\{\mathcal{T}_g\}_{g \in \mathbb{G}}$ and capital \mathcal{T}_k , task prices $\{p(x)\}_{x \in \mathcal{T}}$, task labor demands $\{\ell_g(x)\}_{g \in \mathbb{G}, x \in \mathcal{T}}$ and capital production levels $\{k(x)\}_{x \in \mathcal{T}}$ such that:

E1 Task prices are equal to the minimum unit cost of producing the task:

$$p(x) = \min \left\{ \frac{1}{A_k \psi_k(x)}, \left\{ \frac{w_g}{A_g \psi_g(x)} \right\}_{g \in \mathbb{G}} \right\}.$$

E2 Tasks are produced in a cost-minimizing way, with tasks

$$\mathcal{T}_g = \left\{ x : p(x) = \frac{w_g}{A_g \psi_g(x)} \right\}$$

allocated to workers from skill group g , and tasks

$$\mathcal{T}_k = \left\{ x : p(x) = \frac{1}{A_k \psi_k(x)} \right\}$$

produced with capital.

E3 Task-level employment of labor and capital are given by

$$\ell_g(x) = \begin{cases} y \cdot \frac{1}{M} \cdot A_g^{\lambda-1} \cdot \psi_g(x)^{\lambda-1} \cdot w_g^{-\lambda} & \text{for } x \in \mathcal{T}_g \\ 0 & \text{otherwise.} \end{cases}$$

and

$$k(x) = \begin{cases} y \cdot \frac{1}{M} \cdot A_k^{\lambda-1} \cdot \psi_k(x)^{\lambda-1} & \text{for } x \in \mathcal{T}_k \\ 0 & \text{otherwise.} \end{cases}$$

E4 The labor market clears for all g :

$$\int_{\mathcal{T}_g} \ell_g(x) \cdot dx = \ell_g.$$

E5 The price of the final good is 1, which gives the ideal-price index condition

$$1 = \left(\frac{1}{M} \int_{\mathcal{T}} p(x)^{1-\lambda} \cdot dx \right)^{1/(1-\lambda)}.$$

Fig. 1 provides a graphical illustration of this equilibrium. The task space is represented as a subset of the plane, which is partitioned into $G + 1$ subsets, representing the \mathcal{T}_g 's and \mathcal{T}_k . We explicitly condition these sets on the wage

vector w to emphasize that task allocations depend on wages. The fact that these sets are shown as connected is for simplicity. It can be seen from the figure why the boundaries of these sets, where a task can be produced in a cost-minimizing way by more than one factor, are of measure zero. These sets are determined by comparative advantage, factor-augmenting technologies and factor prices, which influence the costs of performing a task with a given factor.

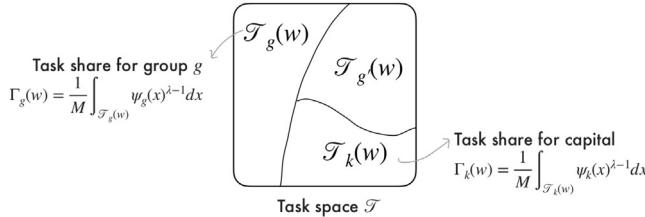


FIGURE 1 Equilibrium task assignment and task shares. The figure depicts the task space and illustrates the assignment of tasks to different groups of workers (g and g' , in this example) and capital (k).

2.3 Equilibrium representation in terms of task shares

Following Acemoglu and Restrepo (2022), we represent and characterize the equilibrium in terms of *task shares*.

Let $\mathcal{T}_g(w)$ be the set of tasks that would be assigned to workers from skill group g at a given level of wages $w = \{w_g\}_{g \in \mathbb{G}}$. Aggregating the labor demand in E3 across tasks, we obtain the labor market-clearing condition

$$\int_{\mathcal{T}_g(w)} y \cdot \frac{1}{M} \cdot A_g^{\lambda-1} \cdot \psi_g(x)^{\lambda-1} \cdot w_g^{-\lambda} \cdot dx = \ell_g.$$

Inverting this equation yields the market-clearing wage for group g ,

$$w_g = \left(\frac{y}{\ell_g} \right)^{1/\lambda} \cdot A_g^{1-1/\lambda} \cdot \Gamma_g(w)^{1/\lambda}, \quad (2)$$

where the task shares are defined as

$$\Gamma_g(w) \equiv \frac{1}{M} \int_{\mathcal{T}_g(w)} \psi_g(x)^{\lambda-1} \cdot dx \quad \text{and} \quad \Gamma_k(w) \equiv \frac{1}{M} \int_{\mathcal{T}_k(w)} \psi_k(x)^{\lambda-1} \cdot dx.$$

Task shares summarize how the market value of tasks assigned to the different groups of workers change as we vary wages. The assumption of strict comparative advantage guarantees that task shares are continuous and differentiable functions of factor prices and technology. Moreover, cost-minimization implies the symmetry property

$$A_{g'}^{1-\lambda} \cdot w_{g'}^{\lambda} \cdot \frac{\partial \Gamma_g(w)}{\partial w_{g'}} = A_g^{1-\lambda} \cdot w_g^{\lambda} \cdot \frac{\partial \Gamma_{g'}(w)}{\partial w_g} \quad \text{for } g' \neq g. \quad (3)$$

This property says that the additional task share that g gains when wages for g' increase equals the additional task share that g' gains when wages for g increase.

Task shares encode all the relevant (local) information on comparative advantage. For example, if the task share of a group decreases by a small (large) amount when its wage increases, this implies that the group has a steep (shallow) comparative advantage at the tasks it currently performs, and cannot be (can be) easily substituted by other groups of workers. Additionally, the behavior of task shares when we increase all wages by the same amount is informative about the substitutability of different groups of workers for capital in marginal tasks.

Proposition 1 (Equilibrium representation). *The competitive equilibrium exists and is unique. The wage vector w and output level y are given by*

$$w_g = \left(\frac{y}{\ell_g} \right)^{1/\lambda} \cdot A_g^{1-1/\lambda} \cdot \Gamma_g(w)^{1/\lambda} \text{ for } g \in \mathbb{G}, \quad (4)$$

$$1 = \underbrace{\left(\Gamma_k(w) \cdot A_k^{\lambda-1} + \sum_g \Gamma_g(w) \cdot \left(\frac{w_g}{A_g} \right)^{1-\lambda} \right)^{1/(1-\lambda)}}_{\equiv \mathcal{C}(w)}, \quad (5)$$

where $\mathcal{C}(w)$ denotes the marginal cost of producing the final good given the wage vector w . The equilibrium level of output can be written as a CES aggregator of the different labor types and capital $k = \int_{\mathcal{T}_k(w)} k(x) dx$, with the equilibrium task shares $\Gamma_g = \Gamma_g(w)$ and $\Gamma_k = \Gamma_k(w)$ appearing as endogenous weights:

$$y = \left(\Gamma_k^{1/\lambda} \cdot (A_k \cdot k)^{1-1/\lambda} + \sum_g \Gamma_g^{1/\lambda} \cdot (A_g \cdot \ell_g)^{1-1/\lambda} \right)^{\lambda/(\lambda-1)}. \quad (6)$$

Like all proofs in this chapter, the proof of this proposition is provided in the Appendix.

Eq. (4) gives the market-clearing wage. This equation demonstrates that equilibrium wages depend on output per worker (y/ℓ_g), factor-augmenting productivity terms (the A_g 's), and the task shares (the $\Gamma_g(w)$'s). Eq. (5) is the ideal-price index condition in E5, rewritten in terms of task shares. This system has a unique solution because task shares satisfy the gross-substitutes property: $\Gamma_g(w)$ is decreasing in w_g and increasing in $w_{g'}$ for all $g' \neq g$.

Eq. (6) is a representation result. Once equilibrium wages and task shares are solved, they can be substituted back into the production function (1) to obtain this form. It shows that the economy behaves as if output were produced using a CES aggregate production function, with the CES weights determined endogenously by equilibrium task shares.

Task shares are the key objects governing the distribution of income in the task model—just as the CES weights govern the distribution of income in a

model with a CES aggregate production function. The share of skill group g in gross national income is⁷:

$$s_g^y = \Gamma_g(w) \cdot \left(\frac{w_g}{A_g} \right)^{1-\lambda}.$$

The share of all labor in gross national income is therefore

$$s_L^y = \sum_g \Gamma_g(w) \cdot \left(\frac{w_g}{A_g} \right)^{1-\lambda} = 1 - \Gamma_k(w) \cdot A_k^{\lambda-1}, \quad (7)$$

and the share of capital in gross national income is

$$s_K^y = \Gamma_k(w) \cdot A_k^{\lambda-1}.$$

Two additional objects of interest are the capital-output ratio, given by

$$\frac{k}{y} = \Gamma_k(w) \cdot A_k^{\lambda-1},$$

and the share of consumption in gross national income, which is

$$\frac{c}{y} = 1 - \Gamma_k(w) \cdot A_k^{\lambda-1}.$$

2.4 Beyond CES

Proposition 1 shows that the task model aggregates to an economy that behaves as if output were produced from a CES aggregator. In this aggregation, task shares determine the resulting CES weights. The fact that task shares are endogenous and depend both on technology and factor prices introduces the two key features that distinguish the task model from previous approaches that rely on CES production functions (or nested versions thereof).

- **Different technologies, different effects:** Technology operates by directly altering the task shares and this enables us to incorporate the distinct impacts of different types of technologies. To see the significance of this feature, suppose we treated (6) as a standard CES production function. Then, the modal form of technology would be a labor-augmenting one, say an increase in A_g , and its effects could be obtained by modifying the first and second terms in the wage equation (4). In this exercise, the elasticity of substitution and the weights would be held constant. In contrast, in our framework, a change in A_g would have a third important effect because it would alter all task shares. More importantly, in the standard framework, we would be forced to think of

⁷ “Gross” here refers to national income inclusive of payments to capital, while net output subtracts these payments.

automation—for example, the introduction of industrial robots—as increasing capital productivity, A_k (this is the only way in which capital can become more productive in that framework). This would have the unambiguous comparative static that it always raises real wages for all worker groups. Instead, in our framework, automation operates entirely by changing task shares and output per worker (the first term), which, as we will see, has very different consequences.⁸

- **Rich substitution patterns:** Despite appearances, the task model does *not* force the elasticity of substitution across groups to equal λ —the elasticity of substitution between tasks. This is because task shares respond to wages, capturing substitution generated by competition for marginal tasks. The task model thus allows for richer substitution patterns than a standard CES model and implies that the resulting macroeconomic elasticities are linked to the pattern of comparative advantage and competition for marginal tasks.

Section 3 introduces a special case of the framework here, which we will refer to as “the no-ripples economy”, to explain the first distinctive feature, while Section 4 discusses the second one and presents a number of simple examples that illustrate the influence of comparative advantage on the macroeconomic elasticity of substitution. Section 5 puts these elements together and characterizes the full implications of different types of technologies in the one-sector model.

3 Different technology, different effects

The first distinctive feature of the task framework is its ability to differentiate between different types of technologies. This section describes the different classes of technology in this model and delineates the distinct mechanisms via which they affect labor demand and productivity. To facilitate the exposition, we focus on a special case of our framework, the “no-ripples economy”, in which there is no competition for marginal tasks.

3.1 The no-ripples economy

We first characterize the impact of different technologies in an example economy that shuts down ripple effects and highlights the distinct direct effects of technology on labor demand. This “no-ripples economy” imposes the following assumption:

Assumption 2 (No ripples). The task space can be partitioned into sets $\{\mathcal{T}_g^*\}_{g \in \mathcal{G}}$ and \mathcal{T}_k^* such that for each g , tasks \mathcal{T}_g^* can be produced only by workers in skill group g and tasks in \mathcal{T}_k^* can be produced only by capital.

⁸ One could try to replicate the effects of automation by exogenously changing the weights of the CES production function, but this has the disadvantage of being highly reduced-form. In particular, there would be no way to know *ex ante* which weights should be changed and by how much.

This assumption ensures that no marginal tasks are being contested between skill groups or between capital and labor. Task shares are pinned down by technology and can be written as

$$\Gamma_g = \frac{1}{M} \int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx \text{ for } g \in \mathbb{G}, \quad \text{and} \quad \Gamma_k = \frac{1}{M} \int_{\mathcal{T}_k^*} \psi_k(x)^{\lambda-1} \cdot dx.$$

From these, one can readily compute equilibrium wages and output using (4) and (6). We maintain Assumption 2 in this section and relax it in subsequent sections.

3.2 Automation

Automation technologies are those that directly displace workers from tasks they perform. In the smartphone production example, automation corresponds to the introduction of robots or computer numerical control machinery that take over various manufacturing and assembly tasks. One can also think of new software systems that automate some of the back-office tasks needed to commercialize smartphones.

We model automation technologies as an increase in the productivity of capital in *tasks previously assigned to labor*. In particular, we assume new automation technologies become available in a set of tasks $\mathcal{A} \subset \cup_{g \in \mathbb{G}} \mathcal{T}_k^*$ and increase capital productivity in these tasks discretely, from $\psi_k(x) = 0$ in $x \in \mathcal{A}$ to $\psi_k^{\text{auto}}(x) > 0$. We assume that in the initial equilibrium $\frac{1}{A_k \cdot \psi_k^{\text{auto}}(x)} < \frac{w_g}{A_g \cdot \psi_g(x)}$ for all $x \in \mathcal{A}$ and for any $g \in \mathbb{G}$. We also assume that \mathcal{A} is a small set (meaning that its measure is small), which guarantees that producing these tasks with capital reduces costs.⁹

A convenient feature of the task framework is that the effects of technology depend on its impact on task allocations and productivity. In the case of automation technologies, we can summarize their effects via two objects: the *direct task displacement* and the *cost savings* that these technologies generate.

Denote the set of tasks performed by skill group g and that now become automated by $\mathcal{A}_g = \mathcal{A} \cap \mathcal{T}_g^*$. The direct task displacement on group g from automating these tasks is

$$d \ln \Gamma_g^{\text{auto}} = \frac{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx} \geq 0.$$

That is, the direct task displacement gives the proportional reduction in group g 's task share resulting from automation—the numerator is the share of tasks in

⁹ Notice that while Assumption 2 holds in the initial equilibrium before the change in technology, it no longer holds after the change, because tasks in \mathcal{A} can be produced by more than one factor of production. The fact that technology changes in a small set of tasks ensures that producing automated tasks with capital is still strictly profitable and thus there are effectively no marginal tasks, even after the change in technology.

the set \mathcal{A}_g , while the denominator is group g 's task share in the initial equilibrium.

The cost savings from automating task x in \mathcal{A}_g are

$$\pi_g^{\text{auto}}(x) = \frac{1}{1-\lambda} \cdot \left(1 - \left[\frac{w_g \cdot A_k \cdot \psi_k^{\text{auto}}(x)}{A_g \cdot \psi_g(x)} \right]^{\lambda-1} \right). \quad (8)$$

This expression measures the decline in costs from switching to produce task x with the new capital instead of labor (at the initial equilibrium wages). Cost savings are positive by assumption. The average cost savings from automating tasks previously assigned to group g can then be computed as the employment-weighted average of $\pi_g^{\text{auto}}(x)$'s:

$$\pi_g^{\text{auto}} = \frac{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot \pi_g^{\text{auto}}(x) \cdot dx}{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx} > 0.$$

Fig. 2 illustrates the role of direct displacement effects from automation and the resulting cost savings for two skill groups.

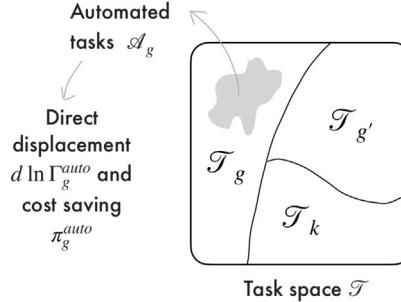


FIGURE 2 Effects of automation on the allocation of tasks. The figure depicts the task space and illustrates an example of new automation technologies increasing the productivity of capital in tasks previously assigned to group g workers, \mathcal{A}_g . This has two consequences: direct displacement and cost savings.

The task displacement and cost-saving gains $\{d \ln \Gamma_g^{\text{auto}}, \pi_g^{\text{auto}}\}_{g \in \mathbb{G}}$ summarize the capabilities of new technologies, the extent to which these capabilities outcompete workers of different skills, and the cost savings generated in the process. The next proposition shows how to compute the effects of automation in terms of these objects.

Proposition 2 (Effects of automation in the no-ripples economy). *The effects of automation technologies, summarized by $\{d \ln \Gamma_g^{\text{auto}}, \pi_g^{\text{auto}}\}_{g \in \mathbb{G}}$, are given by the formulas*

$$d \ln w_g = (1/\lambda) \cdot (d \ln y - d \ln \Gamma_g^{\text{auto}}) \quad \text{for } g \in \mathbb{G} \quad (9)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\sum_g s_g^y \cdot d \ln \Gamma_g^{\text{auto}} \cdot \pi_g^{\text{auto}}}_{=d \ln tfp} . \quad (10)$$

Eq. (9) follows by differentiating (4) and using the fact that task shares are independent of wages in the no-ripples economy. It shows that the impact of automation on wages is given by the sum of two economic forces: the first term, representing the *productivity effect* from automation, and the second term, representing the *displacement effect* from automation—meaning the displacement of workers of group g from the tasks they previously performed. The displacement effect is proportional to $d \ln \Gamma_g^{\text{auto}}$ and is straightforward to compute given the initial equilibrium, as we showed. The productivity effect, on the other hand, depends on how much output increases.

The second equation, (10), which is derived by differentiating (5), can be used to compute the productivity effect and pins down the impact of automation on real wage levels.¹⁰ This equation shows that the average increase in wages equals the TFP gains from automation, which can be computed with a logic identical to Hulten's theorem: $d \ln tfp = \sum_g s_g^y \cdot d \ln w_g$.¹¹

Eq. (10) shows that automation necessarily increases the *average* wage—and does so in proportion to its positive contribution to TFP. The fact that automation increases TFP follows from the fact that, by assumption, capital produces the tasks in \mathcal{A} more cheaply than labor, which implies that $\pi_g^{\text{auto}} > 0$. If this were not the case, these technologies would not be adopted. The result that automation increases average wages in proportion to TFP also has a simple intuition. The change in TFP corresponds to how much the cost of producing the final good declines at given factor prices. Since this cost has to remain at 1, wages must increase on average by some amount proportional to TFP. This result is a consequence of three features: (i) capital is supplied fully elastically (see, for example, Simon, 1965; Caselli and Manning, 2019; Moll et al., 2022; Acemoglu et al., 2024); (ii) all markets are competitive (see Acemoglu and Restrepo, 2024, for the role of labor market imperfections); and (iii) the production technology exhibits constant returns to scale.

The fact that automation increases productivity and *average* wages does not imply that it does so by a significant amount or that it increases *all* workers'

¹⁰ Specifically, the productivity effect $d \ln y$ can be computed by solving Eqs. (9) and (10). This system comprises $G + 1$ unknowns and $G + 1$ equations that can be solved together to determine the changes in the real wage of each group of workers and in output. An equivalent approach uses the fact that $d \ln y = (1 - s_k^y)^{-1} \cdot (d \ln tfp - ds_k^y)$, where s_k^y is the capital share in gross output to obtain this productivity effect and ds_k^y is the change in the capital share, obtained from (11).

¹¹ Hulten's original result focuses on the effects of infinitesimal changes in technology. Here, we have a discrete jump in technology taking place over a small (infinitesimal) set of tasks, but this does not change the overall logic. The only difference is that, when computing $\pi_g^{\text{auto}}(x)$, we have to take into account the impact of this discrete jump on cost shares, which is the reason why the $1 - \lambda$ terms appear in (8).

wages. The formula for the productivity gains from automation shows that these depend on π_g^{auto} . These cost savings can be small—which corresponds to *so-so automation technologies* in Acemoglu and Restrepo (2019). This will be the case when labor is fairly productive in these tasks to start with or when capital can perform these tasks with moderate productivity (just high enough to outcompete labor but not so high as to yield meaningful cost savings). This observation explains why significant investments in automation technologies can generate modest productivity and average wage growth.

Moreover, Eq. (9) highlights that while the productivity effect raises wages on average, the displacement effect can reduce the real wage of affected groups. This can be understood in the simplest way by assuming that automation only affects one group, g , and the new automation technologies are so-so ($\pi_g^{\text{auto}} = \epsilon$ for a small and positive ϵ). One can then show that there is some $\bar{\epsilon}$ such that, for $\epsilon < \bar{\epsilon}$, group g 's real wage will necessarily decline. We return to a detailed discussion of real wage consequences of automation in the presence of ripple effects in Section 5.

The task framework also shows that automation is tightly linked to reductions in the labor share. We can see this from the formula for the capital share in Eq. (7). The expansion in the set of tasks performed by capital implies that the labor share of national income s_L^y decreases—and equivalently, the capital share s_K^y increases—by

$$ds_L^y = - \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{auto}} \cdot (1 + (\lambda - 1) \cdot \pi_g^{\text{auto}}) < 0. \quad (11)$$

This result is a direct consequence of the fact that automation displaces workers from the tasks they used to perform, making production more capital intensive.

Offshoring The task framework can also be used to study the effects of offshoring, which are very similar to automation (see, for example, Grossman and Rossi-Hansberg, 2008). Offshoring corresponds to some tasks previously performed domestically by labor now being transferred to workers in another country. This can be incorporated into our framework by interpreting $k(x)$ to include imports of intermediates (or services) corresponding to task x . For example, the assembly of a smartphone can be performed by robots in the United States, or components can be shipped and assembled in Vietnam. From the viewpoint of workers in the United States, these two shifts have identical effects.¹²

We can therefore model the arrival of new opportunities for offshoring as a jump in the capabilities of the technology used for organizing global supply chains for task x from $\psi_k(x) = 0$ to $\psi_k^{\text{offshore}}(x) > 0$. We define the direct

¹² This is provided that trade is balanced so that a corresponding amount of the final good is transferred to the foreign country to pay for the offshored tasks. In the multi-sector economy studied in the next section, trade balance could be achieved by exporting goods produced in specific industries. If so, the effects of offshoring could differ from automation because they could also involve additional sectoral reallocation.

task displacement from offshoring as $d \ln \Gamma_g^{\text{offshore}}$ and the cost savings from offshoring as π_g^{offshore} analogously as we did for automation.

The objects $\{d \ln \Gamma_g^{\text{offshore}}, \pi_g^{\text{offshore}}\}_{g \in \mathbb{G}}$ summarize the impact of new offshoring opportunities. The effects of offshoring are the same as those in Proposition 2, except that $\{d \ln \Gamma_g^{\text{offshore}}, \pi_g^{\text{offshore}}\}_{g \in \mathbb{G}}$ replace $\{d \ln \Gamma_g^{\text{auto}}, \pi_g^{\text{auto}}\}_{g \in \mathbb{G}}$. The impact of offshoring operates via productivity and displacement effects as well. Just like automation, offshoring can have a negative impact on exposed groups when the cost savings from offshoring are limited.

3.3 New tasks

The second class of technologies considered here are advances that enable the creation of new (labor-intensive) tasks. We emphasized in the Introduction the critical role that new tasks play in generating new opportunities and demand for labor—raising the labor share and counterbalancing the decline in labor share coming from automation. Acemoglu and Restrepo (2018b) and Autor et al. (2024) suggest that a significant part of employment growth over the last six decades is accounted for by occupations in which we see new tasks, such as various technical occupations, radiology, management consulting, design and programming.

While some new tasks emerge as a result of growing preferences for luxury goods (e.g., sommeliers), most new tasks result from advances in technology. For example, radiology became a major occupation because of advances in radiography technology, while management consulting and design occupations are dependent on a range of new communication and design tool innovations. New ride-sharing and delivery jobs were enabled by new platforms leveraging the use of smartphones and GPS technology. Likewise, new consumer products and services often generate new tasks for workers to perform. The defining feature of these examples is that technology creates the demand for new specialized roles or endows workers with new capabilities to produce value and contribute to economic output.

We incorporate new tasks by assuming that there is a technological advance that enables the production of a set \mathcal{N} of new tasks that did not exist in \mathcal{T} . We assume that the sets $\{\mathcal{N}_g\}_{g \in \mathbb{G}}$ have small measure and that, at the initial equilibrium wages, firms strictly prefer to produce tasks in \mathcal{N}_g with workers from skill group g .¹³

The direct effects of new tasks can be summarized by two objects, similar to their counterparts for automation: *direct task reinstatement* and *economic surplus* from new tasks. The direct task reinstatement for group g (driven by the

¹³ We can also allow for new capital-intensive tasks. For example, developing a new design for a widget creates a new task for CNC machinery capable of producing such design. We do not do so to economize space, especially since capital-intensive new tasks do not play as important a role as labor-intensive new tasks in accounting for changes in wage structure.

introduction of new tasks) is

$$d \ln \Gamma_g^{\text{new}} = \frac{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g(w)} \psi_g(x)^{\lambda-1} \cdot dx} \geq 0$$

and gives the percent increase in group g 's task share resulting from the creation of tasks in \mathcal{N} . We refer to this measure as task reinstatement because it corresponds to the expansion of the set of tasks performed by workers in g and is thus the counterpart of the displacement caused by automation.

The economic surplus from new task x in \mathcal{N}_g , evaluated at the initial equilibrium wages, is defined as

$$\pi_g^{\text{new}}(x) = \frac{1}{1-\lambda} \cdot \left(\left[\frac{w_g}{A_g \cdot \psi_g(x)} \right]^{\lambda-1} - 1 \right).$$

The economic surplus from new tasks is positive if the cost of producing the task with labor $w_g/(A_g \psi_g(x))$ is below 1—the price of the final good and our choice of numeraire.¹⁴ We assume this is the case, so that new task x increases TFP and will be adopted. We also define average economic surplus from new tasks for group g as:

$$\bar{\pi}_g^{\text{new}} = \frac{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot \pi_g^{\text{new}}(x) \cdot dx}{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx} > 0.$$

Fig. 3 illustrates the role of direct reinstatement effects from new tasks and the economic surplus this generates.

The objects $\{d \ln \Gamma_g^{\text{new}}, \pi_g^{\text{new}}\}_{g \in \mathbb{G}}$ summarize the reinstatement effect from new tasks and its economic impact. The next proposition shows how to compute the effects of new tasks in terms of these objects.

Proposition 3 (Effects of new tasks in the no-ripples economy). *The effects of new tasks, summarized by $\{d \ln \Gamma_g^{\text{new}}, \pi_g^{\text{new}}\}_{g \in \mathbb{G}}$, are given by the formulas*

$$d \ln w_g = (1/\lambda) \cdot (d \ln y - d \ln M + d \ln \Gamma_g^{\text{new}}) \quad \text{for } g \in \mathbb{G} \quad (12)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\sum_g s_g^y \cdot d \ln \Gamma_g^{\text{new}} \cdot \pi_g^{\text{new}}}_{=d \ln tfp}. \quad (13)$$

¹⁴ Note that new tasks can raise surplus even if $\lambda < 1$. This is because in our framework, the cost function associated with the production of the final good is $c(p) = \left[\frac{1}{M} \int_{\mathcal{T}} p(x)^{1-\lambda} \right]^{1/(1-\lambda)}$. An expansion in the set of tasks can reduce the price index even if $\lambda < 1$, because of the presence of M in the denominator. This modeling approach implies that when new tasks are introduced, the entire production process or organization changes. This is different from the standard way of modeling new varieties, whereby the arrival of a new variety reduces the cost of this latent variety from ∞ to some finite value.

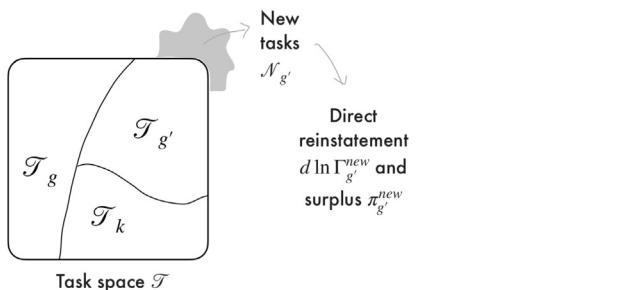


FIGURE 3 Effects of new tasks on the allocation of tasks. This figure depicts the task space and illustrates a change in technology that introduces new tasks, $\mathcal{N}_{g'}$. This has two consequences: direct reinstatement and a surplus.

As in Proposition 2, these two equations can be solved together to determine changes in the real wages of all demographic groups as well as the increase in output. Eq. (12) describes the distributional effects of new tasks. Eq. (13) gives the TFP improvements due to new tasks and pins down their effects on wage levels.

The proposition shows that the wage consequences of new tasks are given by a combination of a *productivity effect* and a *reinstatement effect*, which is the converse of the displacement effect from automation. The reinstatement effect measures the beneficial (positive) impact from new tasks where workers will be employed. In addition, $d \ln M$ is included as a correction term because M , the measure of tasks in the economy, is in the denominator of (1). The assumption that there is a positive economic surplus from new task adoption is sufficient to ensure that average wages increase after accounting for this correction.

Because both the productivity and reinstatement effects are positive, new tasks increase wages for affected groups. Moreover, in contrast to automation technologies, new tasks increase the labor share of national income because they expand the set of tasks performed by labor, making the production process more labor-intensive.¹⁵

3.4 Labor-augmenting technologies

It is useful to distinguish between two types of labor-augmenting technologies. The first (and more realistic) is in the form of new technologies that raise workers' productivity at some of the tasks they currently perform. For example, imagine the creation of a sturdier and lighter hammer, which increases the productivity of male workers without college degrees in construction and carpentry tasks but not in other jobs. We refer to these as *labor-augmenting technology at*

¹⁵ We provide the exact formulas for labor share changes for this and other technologies in the Appendix to save space in the text.

the intensive margin. We represent the effects of labor-augmenting technology at the intensive margin on group g by

$$d \ln \psi_g^{\text{intensive}} = \frac{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot d \ln \psi_g(x) \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx}.$$

This notation emphasizes that these increases in productivity occur only at tasks already assigned to group g .¹⁶

The second alternative involves *uniformly labor-augmenting technological change*, which increases the productivity of a factor in *all tasks* in the economy, and can be represented by increases in the A_g terms. This is the most common type of technological change studied in economic growth models and in previous analyses of inequality. Finding examples of uniformly labor-augmenting technologies is challenging, but one possibility would be assistive technologies that improve the sight of visually impaired workers. The distinction between these labor-augmenting technologies is important in our general framework, though the next proposition shows that in the no-ripples economy, they have identical effects.

Proposition 4 (Labor-augmenting technologies in the no-ripples economy). *The effects of labor-augmenting technologies are given by the formulas*

$$d \ln w_g = (1/\lambda) \cdot \left(d \ln y - (1 - \lambda) \cdot d \ln A_g - (1 - \lambda) \cdot d \ln \psi_g^{\text{intensive}} \right) \quad \text{for } g \in \mathbb{G} \quad (14)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\sum_g s_g^y \cdot (d \ln A_g + d \ln \psi_g^{\text{intensive}})}_{=d \ln \text{tfp}}. \quad (15)$$

Both forms of augmentation affect wages via a productivity effect $d \ln y$. In addition, both forms directly increase worker productivity one-to-one (by $d \ln A_g$ or by $d \ln \psi_g^{\text{intensive}}$), but this has to be weighed against a negative *task-price effect*, given by $(-1/\lambda) \cdot (d \ln A_g + d \ln \psi_g^{\text{intensive}})$. In the no-ripples economy, the task-price effect dominates the quantity expansion for both forms of augmentation in the empirically relevant case where tasks are gross complements ($\lambda < 1$). This means that the benefits from labor-augmenting technologies accrue mostly to other workers who are not themselves becoming more productive and who benefit from the increase in the price of tasks they produce.

That these two forms have identical effects in the no-ripples economy should not be surprising: the set of tasks performed by a factor, say, skill group g , does not change in response to augmenting technologies. Hence a marginal increase

¹⁶ This discussion clarifies that we could alternatively refer to this form of augmentation as “productivity deepening” to capture the fact that it deepens the comparative advantage that the group has for the tasks it is already performing (those in the set \mathcal{T}_g^*).

in A_g only improves the productivity of this factor in the tasks it is performing and is thus very similar to an increase in $d \ln \psi_g^{\text{intensive}}$. For the same reason, labor-augmenting technologies do not affect the labor share of national income in the absence of ripples since none of these technologies alter the range of tasks assigned to capital.¹⁷

It is useful to note the key differences between labor-augmenting technologies and automation and new tasks—a feature that is particularly evident in the no-ripples economy. All of the effects of labor-augmenting technologies are at the intensive margin. They only affect relative wages via task prices, but they do not bring about large changes in the allocation of tasks to factors. In contrast, both automation and new tasks work at the extensive margin—their main impacts are rooted in the changes in the allocation of tasks that they cause. This is also the reason why the balance between the distributional and productivity effects of these types of technologies differ.

To further illustrate this point, we compare the magnitude of the distributional consequences of labor-augmenting and automation technologies (in both cases, relative to their productivity effects). For labor-augmenting technology, this ratio is

$$-\frac{(1 - \lambda) \cdot \psi_g^{\text{intensive}}}{\psi_g^{\text{intensive}}} = -(1 - \lambda).$$

The numerator is the impact via the combination of task price and quantity effects, while the denominator is the increase in their productivity. The corresponding ratio for automation is

$$-\frac{d \ln \Gamma_g^{\text{auto}}}{d \ln \Gamma_g^{\text{auto}} \cdot \pi_g^{\text{auto}}} = -\frac{1}{\pi_g^{\text{auto}}}.$$

The first of these expressions is positive when $\lambda > 1$ (because the quantity effects are larger than the price effects), and even when it is negative, it takes a finite value less than 1. In contrast, the second expression can be unboundedly large, especially for so-so automation technologies ($\pi_g^{\text{auto}} \approx 0$).

Labor-augmenting technologies are also very different from new tasks. While the former increases the quantity of goods and services that workers produce in existing tasks (and this comes at the expense of a reduction in the price of these tasks and services, putting downward pressure on their wages), new tasks reinstate workers into new activities, allowing them to spread their labor across a wider range of tasks. This is the reason why new tasks, which enable the labor hours of the affected group to be distributed across a larger set of tasks, do not run into the same diminishing returns that labor-augmenting improvements do.

¹⁷ This follows from the formula for the labor share in Eq. (7). The equation shows that when the supply of capital is elastic, the labor and capital share are pinned down by the range of tasks assigned to capital and the productivity of capital in these tasks but are independent of labor productivity at other tasks. If the supply of capital were not perfectly elastic, these changes would impact relative task prices and have a (typically small) impact on the labor share.

3.5 Capital-augmenting technologies

The analysis of capital-augmenting changes is similar to that of labor-augmenting ones. For *capital-augmenting technological change at the intensive margin*, we define

$$d \ln \psi_k^{\text{intensive}} = \frac{\int_{T_k} \psi_k(x)^{\lambda-1} \cdot d \ln \psi_k(x) \cdot dx}{\int_{T_k} \psi_k(x)^{\lambda-1} \cdot dx},$$

as the increase in the productivity of capital in the tasks it is already performing. A *uniformly capital-augmenting technological change* is summarized by $d \ln A_k$, analogously to the previous subsection.

Proposition 5 (Capital-augmenting technologies in the no-ripples economy). *The effects of capital-augmenting technologies are given by the formulas*

$$d \ln w_g = (1/\lambda) \cdot d \ln y \quad \text{for } g \in \mathbb{G} \quad (16)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{s_K^y \cdot (d \ln A_k + d \ln \psi_k^{\text{intensive}})}_{=d \ln tfp}. \quad (17)$$

The proposition shows once more the equivalence between intensive-margin and uniformly capital-augmenting technologies in the no-ripples economy. One noteworthy point is that because, in the no-ripples economy, capital-augmenting technologies only change the productivity of already capital-intensive tasks, they do not create any adverse effects on labor, and thus always have a positive impact on the wages of all groups of workers. In fact, when $\lambda < 1$, capital-augmenting technological change at the intensive margin increases the labor share of national income.

This proposition reiterates that there is a crucial difference between capital-augmenting technologies and automation. As already noted, the latter acts exclusively at the extensive margin—by altering the allocation of tasks. Instead, capital-augmenting technologies act primarily (and in the no-ripples economy entirely) at the intensive margin. In fact, while automation reduces the labor share and could reduce the real wage of affected groups, capital-augmenting technologies increase all worker wages uniformly and, in the plausible scenario where capital and labor are gross complements, they also increase the labor share. This distinction clarifies why it would be incorrect to think of the development of industrial robots or other automation technologies as augmenting existing capital.

3.6 Microfoundation for shifting Cobb-Douglas exponents

The no-ripples economy also provides a microfoundation for a Cobb-Douglas aggregate production function where technology acts by changing its elasticities. To see this, consider the limit case with $\lambda \rightarrow 1$. Output in this economy can

then be represented as

$$y = \mathcal{A} \cdot \left(\frac{k}{\Gamma_k} \right)^{\Gamma_k} \prod_g \left(\frac{\ell_g}{\Gamma_g} \right)^{\Gamma_g},$$

where the exponents are given by the share of tasks in \mathcal{T}_g^* and \mathcal{T}_k^* , and $\ln \mathcal{A} = \frac{1}{M} \cdot \int_{x \in \mathcal{T}_k^*} \ln(A_k \cdot \psi_g(x)) \cdot dx + \sum_g \frac{1}{M} \cdot \int_{x \in \mathcal{T}_g^*} \ln(A_g \cdot \psi_g(x)) \cdot dx$.

This example can be used to illustrate several of the conclusions of Propositions 2–5. In particular, we can easily see how automation and new tasks can have sizable effects on the equilibrium by shifting the Cobb-Douglas exponents. In contrast, augmenting technologies work by increasing aggregate productivity \mathcal{A} in a factor-neutral way.

This example also provides a microfoundation for models of factor-eliminating technologies, such as Zuleta (2008) and Peretto and Seater (2013). It shows that one can map automation to a reduction in the Cobb-Douglas exponent for skill groups whose tasks become automated and an increase in the exponent for capital, while new tasks increase the Cobb-Douglas exponent for the favored skill groups and reduce the exponent for capital.

3.7 Taking stock

Several of the key messages discussed in the Introduction are clarified by Propositions 2–5. Most importantly, these results show that new technologies affect equilibrium wages through three mechanisms: a *productivity effect* (any technology that increases productivity and expands output raises labor demand and wages); *displacement and reinstatement effects* (that work at the extensive margin by directly changing the allocation of tasks to factors of production); and *task-price effects* (factor-augmenting technologies increase the supply of some tasks and reduce their prices). Different types of technological changes generate different combinations of these three effects, thus having varied consequences in terms of aggregate productivity and inequality.

4 From micro to macro elasticities

In this section, we focus on the second distinctive feature of the task framework: the rich pattern of macroeconomic elasticities of substitution. This section defines these elasticities and shows how the pattern of comparative advantage shapes them. We then illustrate these patterns with a series of examples.

4.1 Macroeconomic elasticities of substitution

In the no-ripples economy studied in the previous section, any substitution between factors comes only via the substitution between tasks. If high-skill workers become abundant, the tasks they produce also become abundant, driving

down their price and encouraging firms to substitute toward using these tasks more intensively. The general case with ripples allows for richer substitution patterns. As one group of workers becomes abundant, they will also substitute for other workers in marginal tasks. The extent of this effect depends on whether workers compete for marginal tasks and how steep their comparative advantage is in these tasks.

To explore these issues, let us define the macroeconomic elasticity of substitution between skill groups g and g' as

$$\sigma_{gg'} = \frac{1}{s_{g'}^y} \cdot \frac{d \ln \ell_g}{d \ln w_{g'}} \Big|_{y \text{ constant}}.$$

This elasticity measures how much a proportional increase in the wage of skill group g' changes the demand for skill group g . In the task framework, for $g' \neq g$, this elasticity is

$$\sigma_{gg'} = \underbrace{\lambda}_{\text{substitution between tasks}} + \underbrace{\frac{1}{s_{g'}^y} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}}}_{\text{substitution within marginal tasks}}.$$

With constant returns to scale, the elasticity is symmetric: $\sigma_{gg'} = \sigma_{g'g}$.¹⁸

The formula illustrates the two margins of substitution. First, we have substitution between tasks produced by different skill groups and controlled by λ . This is similar to the substitution in the standard CES production function and is the only margin of substitution in the no-ripples economy. Second, we have substitution between worker groups taking place in marginal tasks. This second source of substitution depends on the intensity of competition for marginal tasks and is shaped by the comparative advantage schedules. This term will be high when the two groups in question have similar comparative advantage schedules in marginal tasks, which in turn would imply that a small difference in costs of producing these marginal tasks can lead to a big shift in tasks from one group to the other.¹⁹

¹⁸ The notion of elasticity of substitution used here is due to Allen-Uzawa. With constant returns to scale, the Allen-Uzawa elasticity can be expressed in terms of the cost function $\mathcal{C}(w)$ as

$$\sigma_{gg'} = \frac{\mathcal{C}(w) \cdot \mathcal{C}_{gg'}(w)}{\mathcal{C}_g(w) \cdot \mathcal{C}_{g'}(w)},$$

which is symmetric due to Young's theorem. Note that the symmetry of $\sigma_{gg'}$ is equivalent to the symmetry property in (3), also proving that assertion.

¹⁹ Macroeconomic elasticities of substitution can be estimated from the data, but the exact source of variation being exploited is important. If one focuses on situations in which tasks cannot be or are not reassigned between factors of production, then one would recover λ .

The elasticity of substitution between capital and skill group g can be similarly computed as:

$$\sigma_{kg} = \frac{1}{s_g^y} \cdot \frac{\partial \ln k}{\partial \ln w_g} \Big|_{y \text{ constant}} = \underbrace{\lambda}_{\text{substitution between tasks}} + \underbrace{\frac{1}{s_g^y} \frac{\partial \ln \Gamma_k(w)}{\partial \ln w_g}}_{\text{substitution within marginal tasks}}.$$

The two margins of substitution are present in this case as well and play a central role in determining how advances in the productivity of capital in marginal tasks impacts workers (see Acemoglu and Loebbing, 2024).

4.2 Examples

This subsection illustrates how the macroeconomic elasticity of substitution is determined in a number of tractable cases, clarifying the role of comparative advantage.

Equilibrium with a common elasticity of substitution between tasks The simplest example of how the macroeconomic elasticity of substitution is determined by the pattern of comparative advantage comes from Acemoglu and Zilibotti (2001), who analyze a task model with two types of labor: low-skill (with supply ℓ) and high-skill (with supply h). The task space is a line from $[0, 1]$ (so that $M = 1$), tasks are combined with an elasticity of substitution $\lambda = 1$, and

$$y(x) = A_\ell \cdot (1 - x)^{1/\kappa} \cdot \ell(x) + A_h \cdot x^{1/\kappa} \cdot h(x), \text{ where } \kappa > 0.$$

In this economy, task shares can be computed as

$$\Gamma_\ell(w) = \frac{(w_h/A_h)^\kappa}{(w_h/A_h)^\kappa + (w_\ell/A_\ell)^\kappa}, \quad \Gamma_h(w) = \frac{(w_\ell/A_\ell)^\kappa}{(w_h/A_h)^\kappa + (w_\ell/A_\ell)^\kappa},$$

and the macroeconomic elasticity of substitution between low and high-skill labor is constant and given by

$$\sigma_{h\ell} = \underbrace{\frac{1}{s_\ell^y}}_{=\lambda} + \frac{1}{s_\ell^y} \cdot \frac{\partial \ln \Gamma_h(w)}{\partial \ln w_\ell} = 1 + \frac{1}{s_\ell^y} \cdot (1 - s_h^y) \cdot \kappa = 1 + \kappa.$$

In fact, the equilibrium admits a representation that takes the following CES form:

$$y = \left((A_\ell \cdot \ell)^{\frac{\kappa}{1+\kappa}} + (A_h \cdot h)^{\frac{\kappa}{1+\kappa}} \right)^{\frac{1+\kappa}{\kappa}}.$$

We see in this example that the macroeconomic elasticity of substitution between low and high-skill is $1 + \kappa > 1$, different both from the (infinite) within-task elasticity of substitution and the elasticity of substitution between tasks

(which is equal to $\lambda = 1$). Intuitively, a greater value for κ makes the comparative advantage of high-skill labor relative to low-skill labor shallower in marginal tasks, facilitating the assignment of more tasks to the type of labor that is cheaper. In contrast, when κ is low, the productivity of high-skill labor relative to low-skill labor declines sharply as more tasks are assigned to high-skill workers.

Macroeconomic elasticity of substitution with correlated Frechet distributions

This example generalizes the previous one to a setting with multiple (> 2) skill groups. It is also an adaptation of the commonly-used parameterization of Eaton and Kortum (2002) of the original Dornbusch et al. (1977) model, with skill groups taking the place of countries and no trade costs.²⁰ This example illustrates how correlation and (lack of dispersion) in task-level productivities makes skill groups more substitutable in the aggregate.

Consider a version of the task model with multiple types of workers and no capital. The task space is a line from $[0, 1]$ (so that $M = 1$), tasks are combined with an elasticity of substitution $\lambda \in (0, 1)$, and

$$y(x) = \sum_g A_g \cdot \psi_g(x) \cdot \ell_g(x).$$

Suppose that for each task x , the task-level productivities of the different worker groups $\psi_g(x)$ are drawn from a correlated Frechet distribution with CDF:

$$\Pr(\psi_1(x) \leq a_1, \dots, \psi_G(x) \leq a_G) = \exp \left\{ - \left[\sum_g a_g^{-\kappa/(1-\rho)} \right]^{1-\rho} \right\}.$$

In this specification, $\rho \in [0, 1]$ measures the correlation between the productivities of different groups of workers, and $\kappa > 0$ is an inverse measure of dispersion in productivities. The case $\rho = 0$ gives the commonly used case of independent Frechet distributions.

In this example, task shares can be computed as

$$\Gamma_g(w) = \left(\frac{w_g}{A_g} \right)^{\lambda-1-\kappa/(1-\rho)} \cdot \left[\sum_{g'} \left(\frac{w_{g'}}{A_{g'}} \right)^{-\kappa/(1-\rho)} \right]^{\frac{\lambda-1-\kappa/(1-\rho)}{\kappa/(1-\rho)}},$$

which implies a common macroeconomic elasticity of substitution between skill groups

$$\sigma_{gg'} = \underbrace{\lambda}_{\text{between tasks}} + \underbrace{\frac{1}{s_{g'}^y} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}}}_{\text{within marginal tasks}} = \lambda + \left(\frac{\kappa}{1-\rho} - \lambda + 1 \right).$$

²⁰ Lind and Ramondo (2023) utilize this parametrization in a trade context, while Dvorkin and Monge-Naranjo (2019) and Freund (2024) use it in task models.

Equilibrium output again aggregates to a CES representation, this time with elasticity $1 + \kappa/(1 - \rho)$ and productivity level \mathcal{A} (for some constant \mathcal{A}):

$$y = \mathcal{A} \cdot \left(\sum_g (A_g \cdot \ell_g)^{\frac{\kappa}{1-\rho+\kappa}} \right)^{\frac{1-\rho+\kappa}{\kappa}}.$$

The macroeconomic elasticity of substitution, $1 + \kappa/(1 - \rho)$, exceeds λ because it accounts for substitution in marginal tasks. Note that when κ is larger, skills are less dispersed, and comparative advantage across workers is shallower, translating into greater substitution between worker types. Substitution in marginal tasks also increases with ρ . Greater correlation in workers' productivity implies a more intense competition for marginal tasks.

The macroeconomic elasticity of substitution between capital and labor The setup of Hubmer and Restrepo (2021) provides an example where tasks are complements but the macroeconomic elasticity of substitution between capital and labor becomes 1.

Suppose that there are two factors of production: labor ℓ and capital k . The task space is the line $[0, 1]$ (so that $M = 1$) and tasks are combined with an elasticity $\lambda \in (0, 1)$. Suppose also that the productivities of capital and labor in task x are

$$\psi_k(x) = x^{\frac{1-1/\gamma_k}{1-\lambda}} \cdot (1-x)^{\frac{1+1/\gamma_k}{1-\lambda}} \quad \text{and} \quad \psi_\ell(x) = x^{\frac{1+1/\gamma_\ell}{1-\lambda}} \cdot (1-x)^{\frac{1-1/\gamma_\ell}{1-\lambda}}.$$

Equilibrium output now takes a Cobb-Douglas form

$$y = \mathcal{A} \cdot k^{\frac{\gamma_k}{\gamma_k+\gamma_\ell}} \cdot \ell^{\frac{\gamma_\ell}{\gamma_k+\gamma_\ell}}$$

and we can also see that the macroeconomic elasticity of substitution between capital and labor is unity. This is because, in this case, the additional substitution coming from the comparative advantage schedules adds to the elasticity of substitution between tasks, $\lambda < 1$. The γ parameters determine the importance of capital and labor in this Cobb-Douglas aggregator.

5 Putting it all together: shocks and propagation in the one-sector economy

In this section, we provide a characterization of the full equilibrium in the one-sector economy, bringing together the analysis of different types of technologies from Section 3 and the macroeconomic patterns of substitution from Section 4. The main tool for this analysis is the propagation matrix, which we introduce in the next subsection. We will also see that the effects of different types of technologies are richer in this case because of the substitution patterns that they initiate. Throughout, we focus on first-order approximations to the equilibrium effects of various changes, meaning that the formulas we present apply to small changes in technology.

5.1 Equilibrium: ripple effects and the propagation matrix

In the no-ripples economy, technology affected task shares directly. For example, in Proposition 2 automation reduces exposed groups' relative wage and potentially their real wage via a displacement effect. More generally, however, once group g experiences a decline in its relative wage, it becomes more profitable for some firms to use this group of workers in marginal tasks, substituting for other groups and putting pressure on their wages. This competition for marginal tasks is the source of *ripple effects*, which capture the indirect consequences of the reallocation of tasks between groups.

Fig. 4 illustrates the role of ripple effects in an example where automation displaces workers from group g and new tasks are created for group g' . Both technological developments increase the relative wage of group g' , encouraging firms to substitute capital or workers from skill group g for those from group g' in marginal tasks. This endogenous reallocation of tasks is depicted by the dotted lines.

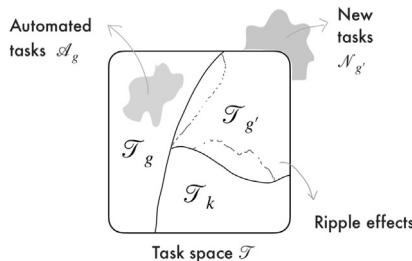


FIGURE 4 Direct effects of technology and ripple effects. The figure depicts the task space and shows the direct and the ripple effects caused by automation and new tasks (dotted and dashed lines).

To understand the implications of ripple effects, consider a demand shock affecting group g . This could be automation, labor-augmenting technological change, new tasks or other forms of technology. In the no-ripples economy, the impact of this shock on group g can be decomposed into its productivity $d \ln y$ and direct effects z_g , so that $d \ln w_g = (1/\lambda) \cdot (d \ln y + z_g)$. In the general case with ripples, differentiating the wage equation (4) yields

$$d \ln w_g = \frac{1}{\lambda} \cdot d \ln y + \frac{1}{\lambda} \cdot z_g + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \cdot d \ln w, \quad (18)$$

where $d \ln w = (d \ln w_1, \dots, d \ln w_G)$ is the column vector of all wage changes. These wage changes affect the equilibrium wage of group g by reallocating marginal tasks. This effect is summarized by the Jacobian $\partial \ln \Gamma_g(w)/\partial \ln w$, written as the row vector of marginal changes in group g 's task share.

Stacking (18) for all groups and collecting the terms involving $d \ln w$ on the left-hand side allows us to solve for the endogenous change in wages as a function of the vector (z_1, z_2, \dots, z_G) . In what follows, we use the notation $\text{stack}(z_g)$ to represent this vector.

Proposition 6 (Effects of technology with ripple effects). *Consider a set of technological changes with direct effects $\text{stack}(z_g)$, which jointly reduce the marginal cost of producing the final good by $\pi = -d \ln \mathcal{C}(w)|_{w=\text{constant}} > 0$ holding all wages constant. The effect of these technological changes on wages and output is given by*

$$d \ln w = \Theta \cdot \text{stack}(d \ln y + z_g) \quad (19)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\pi}_{=d \ln tfp}, \quad (20)$$

where

$$\Theta = \frac{1}{\lambda} \cdot \left(1 - \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \right)^{-1}$$

is the propagation matrix.

Eq. (19) provides a general formula that applies to all forms of technological change. It shows that we can decompose the effects of any technology into a productivity effect $d \ln y$, direct effects z_g (which include task displacement, task reinstatement, and task-price substitution effects), and the ripple effects subsumed in the propagation matrix Θ . The reason why the propagation matrix takes the form of a Leontief inverse is that it accumulates the impacts resulting from the reallocation of marginal tasks between g and g' , which then leads to a second round of reallocation of marginal tasks between g' and g'' , and so on. Eq. (20), on the other hand, shows that the TFP gains and average wage increase due to technology are the same as in the no-ripples case. This is because our economy is competitive and, with the standard envelope theorem logic, the substitution of one group of workers for another at marginal tasks does not generate any first-order gains in productivity. These $G + 1$ equations can again be solved together to obtain wage changes for the G groups of workers and the change in output for the unique final good.

5.2 Properties of the propagation matrix

When there is no competition for marginal tasks, as in the case studied in the no-ripples economy, the propagation simplifies to

$$\Theta = \begin{pmatrix} \frac{1}{\lambda} & 0 & \dots & 0 \\ 0 & \frac{1}{\lambda} & \dots & 0 \\ & & \ddots & \\ 0 & 0 & \dots & \frac{1}{\lambda} \end{pmatrix},$$

and we recover the formulas for the no-ripples economy.

For the general case, the Appendix establishes that the propagation matrix is well defined and has non-negative entries. The off-diagonal entries $\theta_{gg'} \geq 0$ capture the extent to which group g' competes directly or indirectly (via subsequent rounds of reassignment) for marginal tasks with workers in group g .

The propagation matrix has several important properties:

- Dampening:** All eigenvalues of Θ are real and in the $[0, 1/\lambda]$ interval. This means that ripple effects dampen the distributional consequences of a shock. Intuitively, once a group is able to compete for and take over marginal tasks from others, the burden of the direct shocks it suffers will be lessened. This force exhibits itself by the diagonal element of Θ corresponding to group g being less than $1/\lambda$ (recalling that the direct effect of a shock is $(1/\lambda) \cdot z_g$).
- Monotonicity:** for all $g' \neq g$, we have

$$\theta_{gg} \geq \theta_{g'g},$$

so that the maximum entry along a column of the propagation matrix is in the diagonal. This implies that a shock directly increasing (reducing) demand for g cannot increase (decrease) the wage of group g' by more than g 's wage. This monotonicity property ensures that relative demand curves for skill groups are downward sloping.

- Row sums:** Row sums of the propagation matrix are

$$\rho_g = \sum_{g'} \theta_{gg'} = \frac{1}{\lambda} \cdot \left[1 + s_K^y \cdot \left(\frac{\bar{\sigma}_{kg}}{\lambda} - 1 \right) \right]^{-1} \quad \text{for } g \in \mathbb{G},$$

where $\bar{\sigma}_{kg} = \sum_{g'} (\theta_{gg'} / \rho_g) \cdot \sigma_{kg'}$ and s_K^y is the share of capital in national income. In the special case where there is no capital, this simplifies to $\rho_g = \sum_{g'} \theta_{gg'} = 1/\lambda$ for all groups. Another noteworthy special case is when all groups are equally substitutable with capital, i.e., $\sigma_{kg} = \sigma_k$, in which case we have

$$\rho_g = \sum_{g'} \theta_{gg'} = \frac{1}{\lambda} \cdot \left[1 + s_K^y \cdot \left(\frac{\sigma_k}{\lambda} - 1 \right) \right]^{-1} \quad \text{for } g \in \mathbb{G}.$$

The comparison of these two expressions shows that skill groups that are more substitutable for capital tend to have lower row sums.

- Propagation and substitution:** The propagation matrix Θ is related to the matrix of elasticities of substitution $\Sigma = \{\sigma_{gg'}\}_{g,g' \in \mathbb{G}}$ via the identity

$$\Theta = \text{diag} \left(\frac{1}{s^y} \right) \cdot (\lambda - \Sigma)^{-1},$$

where $\text{diag}(1/s^y)$ is a diagonal matrix with entries $(1/s_1^y, \dots, 1/s_G^y)$. This equation thus clarifies the tight connection between ripple effects and substitutability between labor types—greater substitution generates more substantial ripple effects and leads to smaller diagonals in the propagation matrix.

5. Symmetry: The propagation matrix satisfies the symmetry property $\theta_{gg'}/s_{g'}^y = \theta_{g'g}/s_g^y$ —a corollary of the symmetry of task shares and elasticities of substitution.

To illustrate these properties, we can return to the examples introduced above. In the Frechet example, the propagation matrix is

$$\Theta = \begin{pmatrix} \frac{1}{\kappa/(1-\rho)+1} + \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_1^y & \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_2^y & \dots & \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_G^y \\ \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_1^y & \frac{1}{\kappa/(1-\rho)+1} + \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_2^y & \dots & \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_G^y \\ \dots & \dots & \dots & \dots \\ \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_1^y & \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_2^y & \dots & \frac{1}{\kappa/(1-\rho)+1} + \frac{\kappa/(1-\rho)+1-\lambda}{(\kappa/(1-\rho)+1)\cdot\lambda} \cdot s_G^y \end{pmatrix}.$$

With the Frechet parameterization, ripple effects are uniform—so that a shock to group g creates the same wage consequences across all other groups. All eigenvalues of this matrix are equal to $1/(\kappa/(1-\rho)+1)$, and thus all shocks are damped by $\lambda/(\kappa/(1-\rho)+1)$. Naturally, the task framework is more general and allows for richer (and less restrictive) propagation patterns.

In the rest of this section, we study how different types of technological and factor supply changes impact the economy via their direct effects and their indirect effects working through the propagation matrix.

5.3 Automation

We first use Proposition 6 to study the implications of automation technologies (as in Section 3, the same results apply to offshoring and we do not repeat those here).

Consider new technologies leading to the automation of the set of tasks $\mathcal{A} = \cup_g \mathcal{A}_g$ (with the same convention as before that \mathcal{A}_g comprises tasks previously performed by skill group g). Let us also assume, for simplicity, that, for each g , \mathcal{A}_g is in the interior of the set of tasks performed by this group, \mathcal{T}_g . Then we can again summarize the share of tasks lost to automation for each skill group by $\{d \ln \Gamma_g^{\text{auto}}\}_g$, and cost savings from automation can be written as $\pi = \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{auto}} \cdot \pi_g^{\text{auto}}$, where π_g^{auto} is the average cost savings from automating tasks previously performed by skill group g .

Proposition 6 implies that the effects of automation on wages are given by

$$d \ln w = \Theta \cdot \text{stack} \left(d \ln y - d \ln \Gamma_g^{\text{auto}} \right) \quad (21)$$

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\sum_g s_g^y \cdot d \ln \Gamma_g^{\text{auto}} \cdot \pi_g^{\text{auto}}}_{=d \ln \text{tfp}}. \quad (22)$$

Eq. (9) from the no-ripples economy is a special case of (21), with the propagation matrix replaced by a matrix with $1/\lambda$ on the diagonal. All discussion of that equation applies in this case as well: automation again works via the productivity effect summarized by the increase in output and the displacement effects summarized by $d \ln \Gamma_g^{\text{auto}}$.

Importantly, however, the full distributional effects of automation differ from those in the special case with no ripples. In the general case, groups of workers displaced from their tasks by automation intensify the competition for marginal tasks against groups with whom they are highly substitutable. This competition mitigates the adverse effects of automation on exposed groups by spreading the incidence of this shock more broadly. The formula for wages in (21) shows that, in equilibrium, the downward wage pressure exerted by automation on a group not only depends on the displacement it experiences directly, as in the no ripple case, but also on whether groups competing for marginal tasks are being displaced, and groups competing against these groups are being displaced, and so on, as accounted for by the propagation matrix.

The TFP impact of automation in Eq. (22) is identical between the economies with and without ripples. This is because of the same envelope theorem logic explained above. The reason why the automation shock itself has an impact on TFP is that it is not second-order—it corresponds to a discrete increase in the productivity of capital in a small set of tasks.

When does automation reduce real wages? We can use the general formula for the wage effects of automation in Eq. (21) to identify the circumstances that can lead to real wage declines for exposed groups of workers.

As we have seen, the combination of competitive markets, constant returns to scale production possibilities, and a fully elastic supply of capital ensure that automation increases real wages on average. This is true in the economy with ripples as well as in the no-ripples economy. However, this positive average wage effect can coexist with significant negative impacts on some groups of workers. Proposition 6 allows for a sharper characterization of the conditions under which negative wage effects can arise.

From Eq. (21), the full impact of automation technologies on group g is

$$d \ln w_g = \rho_g \cdot d \ln y - \sum_{g'} \theta_{gg'} \cdot d \ln \Gamma_{g'}^{\text{auto}},$$

where ρ_g is the g th row sum of the entries of Θ .

Three conditions are needed for automation to reduce the real wages of group g :

- i the task displacement from automation concentrates on group g ;
- ii group g is not highly substitutable with unaffected groups of workers;
- iii the cost savings from automation are limited, or automation is “so-so.”

The example outlined in our discussion of Proposition 2 satisfies these three requirements. In the example, we consider a case in which $d \ln \Gamma_g^{\text{auto}} > 0$ and $d \ln \Gamma_{g'}^{\text{auto}} = 0$ for all other groups. This means that the displacement effect of automation is highly concentrated on group g as opposed to being equally shared among all workers. The example was also given in the context of the no-ripples economy. Because there are no marginal tasks in this economy, exposed groups

(and, in fact, all groups) are not highly substitutable. As a result, the propagation matrix is diagonal, with entries $1/\lambda$, and all groups bear the full incidence of any labor demand shock affecting them. Finally, and as discussed above, this form of automation reduces the wages of a group g when $\pi_g^{\text{auto}} = \epsilon$ for some positive ϵ smaller than $\bar{\epsilon}$.

To understand why these three conditions are needed, let us modify our example. Imagine first that the task displacement from automation is not concentrated on a handful of groups and suppose, on the contrary, that automation is fully even across groups, that is, $d \ln \Gamma_g^{\text{auto}} = d \ln \Gamma^{\text{auto}} > 0$. Proposition 2 implies that in the case with no ripples $d \ln w_g = (1/\lambda) \cdot (d \ln y - d \ln \Gamma^{\text{auto}})$ and Proposition 6 shows that in the case with ripples this extends to $d \ln w_g = \rho_g \cdot (d \ln y - d \ln \Gamma^{\text{auto}})$. In both cases, wages change by a proportional amount. This makes intuitive sense, as all workers share the productivity gains and displacement effects from automation evenly. From the fact that average wages must increase following any technological advance, we can conclude that $d \ln y - d \ln \Gamma^{\text{auto}} > 0$ and no group can experience a real wage decline.

Next, to understand the role of substitutability of different work groups, consider the polar opposite of the no-ripples economy, where task-level productivities are highly correlated across groups. In this case, worker groups compete strongly for marginal tasks and become highly substitutable in the aggregate. For example, consider the economy with correlated Frechet draws discussed in the previous section and focus on the limit case where the correlation parameter ρ goes to 1. In the limit, the propagation matrix converges to

$$\Theta = \begin{pmatrix} s_1^y & s_2^y & \dots & s_G^y \\ s_1^y & s_2^y & \dots & s_G^y \\ \vdots & \vdots & \ddots & \vdots \\ s_1^y & s_2^y & \dots & s_G^y \end{pmatrix}.$$

Proposition 6 implies that in this case all wages change by an equal amount $d \ln w_g = d \ln y - \sum_{g'} s_{g'}^y d \ln \Gamma_{g'}^{\text{auto}} > 0$ (and this holds even if the displacement effects from automation were uneven to begin with). Intuitively, when workers compete very strongly for marginal tasks, ripple effects will be equal to direct effects, and the incidence of a demand shock is evenly shared across all workers. Then, because average wages increase following any technological advance, all groups must experience a common real wage increase.

The role of cost savings was discussed in detail above, and large cost savings imply that, regardless of the presence of ripple effects, the productivity gains dominate the displacement effect for all groups, leading to an increase in real wages for all. This reasoning establishes that only “so-so” automation technologies can reduce the wages of exposed workers.

One way to summarize this discussion is as follows. Automation has two effects: it raises group wages on average and creates dispersion around that common wage increase. The common level shift depends on how sizable the cost

savings from automation are. The dispersion or inequality brought by automation depends on how concentrated the shock is and the extent to which workers bear or spread the incidence of this shock. If the shock is evenly spread or the incidence is widely shared (because workers are highly substitutable in marginal tasks), automation will have limited effects on inequality, and all groups will see their real wages increase. Otherwise, automation will have sizable effects on inequality. The cost savings will then determine whether workers who lost in relative terms from automation will also lose in real terms.

To conclude our discussion, we note that automation can also reduce the wages of groups that are not directly exposed to it but are highly substitutable with exposed groups. For example, imagine two groups whose task-level productivities are highly correlated. In the limit, these groups have identical rows in the propagation matrix. Proposition 6 then implies that if automation reduces the real wage of one of these groups, it must also reduce the wages of the other via their strong competition for marginal tasks. This example explains why automating tasks held by middle-skilled workers can also reduce wages at the bottom of the wage distribution.

5.4 New tasks

Proposition 6 generalizes Proposition 3 in the case of new tasks. The full effects of new tasks on wages and output are now given by

$$\begin{aligned} d \ln w &= \Theta \cdot \text{stack} \left(d \ln y - d \ln M + d \ln \Gamma_g^{\text{new}} \right) \\ d \ln tfp &= \sum_g s_g^y \cdot d \ln w_g = \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{new}} \cdot \pi_g^{\text{new}}. \end{aligned}$$

Wages depend on a productivity effect, a task reinstatement effect, and ripples, which account for the propagation of shocks across worker groups due to the endogenous reassignment of tasks. Note that here, ripple effects generate a positive impact on other groups, even if they do not benefit from new tasks directly. This is because workers who obtain new tasks become more expensive and thus less competitive for previously marginal tasks, increasing the demand for other skill groups in those tasks.

5.5 Labor-augmenting technology

In the presence of ripple effects, uniformly labor-augmenting technologies and increases in productivity at the intensive margin have different impacts. The results from Proposition 6 extend to these technologies and imply that their effects on wages are now given by

$$\begin{aligned}
d \ln w = & \Theta \cdot \underbrace{(d \ln y - (1 - \lambda) \cdot \text{stack}(d \ln A_g + d \ln \psi_g^{\text{intensive}}))}_{\text{negative task-price decline from no-ripples case}} \\
& + \underbrace{(1 - \Theta\lambda) \cdot \text{stack}(d \ln A_g)}_{\text{reallocation from uniform improvements}}, \tag{23}
\end{aligned}$$

while the contribution of these technologies to TFP (which pins their effect on wage levels) is the same as in the no-ripples economy.²¹

Labor-augmenting technologies at the intensive margin affect wages via a productivity effect and via the same adverse task-price declines we saw for the no-ripples economy. These effects then propagate via Θ .

Uniform labor-augmenting technologies additionally allow groups becoming more productive to outcompete others for marginal tasks, increasing their task shares. This reallocation is also governed by the propagation matrix, which explains the extra term $(1 - \Theta\lambda) \cdot \text{stack}(d \ln A_g)$ in the equation. This is always beneficial for own wages because $1 - \Theta\lambda$ has a positive diagonal (and also negative off-diagonals, which correspond to marginal tasks being lost to other groups that have become more productive). This positive benefit dominates the adverse price declines at the intensive margin if θ_{gg} is below one, meaning that, group g has a sufficiently high macroeconomic elasticity of substitution with other skill groups.

This discussion further clarifies the difference between (uniform) factor-augmenting technological change—the form of technological progress typically emphasized in the literature on skill-biased technical change building on Katz and Murphy (1992)—and automation, as analyzed in Acemoglu and Restrepo (2022). In particular, Eq. (23) clarifies that the distributional effects of factor-augmenting improvements in technology are fully mediated by the macroeconomic elasticities of substitution, summarized by the propagation matrix. If macroeconomic elasticities are not far from unity, as many available estimates suggest, factor-augmenting technologies will have modest distributional effects. Put differently, with macroeconomic elasticities close to unity, one would need very large increases in group-level productivities to generate a meaningful divergence in wages across groups. In contrast, automation works at the extensive margin, and if it displaces low-education groups from tasks they were previously performing, its direct impacts could be much larger—regardless of the macroeconomic elasticities of substitution since its main effect work by directly changing task shares. This explains why automation can have sizable distributional consequences, even when different factors of production have macroeconomic

²¹ In contrast to the no-ripple economy, labor-augmenting technologies can now change the labor share, and whether the labor share increases or decreases depends on how the task share of capital changes.

elasticities of substitution near one.²² We return to this issue in Section 8, where we explore this distinction quantitatively (see also the discussion in Acemoglu and Restrepo, 2020b).

5.6 Capital-augmenting technologies

In an economy with ripples, capital-augmenting technological change at the intensive margin and uniformly capital-augmenting technological change have different implications. The results from Proposition 6 extend to these technologies and imply that their effects on wages are now given by

$$d \ln w_g = \rho_g \cdot d \ln y - (1 - \lambda \cdot \rho_g) \cdot d \ln A_k$$

while the effects on TFP are identical to those in the no-ripples economy. In this expression, $\rho_g \in [0, 1/\lambda]$ are the row sums of the propagation matrix. As in Proposition 5, capital-augmenting technologies at the intensive margin benefit all worker groups because they make capital more productive, generating a productivity effect, but they do not make capital more competitive in any marginal tasks. In contrast, the implications of uniformly capital-augmenting technologies differ because they now make capital more competitive in marginal tasks. This extra competition is captured by the negative term $(1 - \lambda \cdot \rho_g) \cdot d \ln A_k$, where a larger difference between 1 and $\lambda \cdot \rho_g$ signifies that group g is more substitutable for capital in marginal tasks.

As with uniform labor augmenting technologies, we see here that the distributional effects of uniform capital augmenting technologies are entirely determined by the macro elasticities of substitution between capital and labor, which are subsumed in the row sums of the propagation matrix. If these elasticities are not far from unity, uniform advances in capital, as those considered in Krusell et al. (2000) and the literature on investment-specific technical change, do not generate sizable distributional effects. Moreover, if these macro elasticities are below one, uniform advances in capital cannot generate the observed decline of the labor share. This contrasts with our findings for automation. The effects of automation on the wage distribution and factor shares are fully decoupled from these macro elasticities because automation shifts the allocation of tasks directly at the extensive margin.

The formulas above provide a different microfoundation for skill-specific elasticities of substitution between capital and labor (a possibility first considered by Griliches, 1969). As an example, consider an economy with two types of labor, low-skill and high-skill. Suppose that high-skill labor has a very steep

²² A related distinction explained in Acemoglu (2002) and Acemoglu and Autor (2011) is that, in canonical models of skill-biased technical change with two skill groups, technological change that makes highly-educated workers more productive necessarily increases wages for the low-education group (an implication of q -complementarity with two factors of production and constant returns to scale). Instead, and as shown here, models of automation can generate large wage declines for exposed groups.

comparative advantage schedule in tasks in which it is competing against capital, while low-skill labor has a flatter comparative advantage. A uniform increase in capital-augmenting technological change will then increase inequality because it *de facto* complements high-skill labor, while creating a more intense competition against low-skill labor.

5.7 Changes in labor supply

The following proposition shows that the propagation matrix also mediates the effect of labor supply changes.

Proposition 7 (Effects of exogenous changes in labor supply). *The effects of exogenous changes in $\{\ell_g\}_{g \in \mathbb{G}}$ are given by*

$$d \ln w = \Theta \cdot \text{stack}(d \ln y - d \ln \ell_g) \quad (24)$$

where $d \ln y$ is pinned down by $\sum_g s_g^y \cdot d \ln w_g = 0$.

Labor supply changes affect the wage structure through the propagation matrix because a labor supply expansion generates competition for marginal tasks from the expanding groups. This competition then determines the impact on the wages of both the expanding group and others. The propagation matrix summarizes these cross-group elasticities as well as the demand elasticity for the affected group. The substitution patterns summarized in the propagation matrix also point to the possibility that a particular group (say, domestic low-education workers) may suffer lower wages because of the increase in the supply of another group that is highly substitutable to them (such as immigrant workers).²³

This proposition also provides guidance on how to account for the effects of exogenous labor supply changes on the wage structure, generalizing the approach in Katz and Murphy (1992) and Card and Lemieux (2001), who assume that substitution patterns are given by a nested CES.

²³ In this case, we would have that the two groups are q -substitutes (as opposed to the more standard notion of q -complementarity). The propagation matrix contains all relevant information on whether different skill groups are q -complements or q -substitutes. Consider, for example, a case with no capital. An increase in the supply of skill group g increases output by $d \ln y = s_g^y \cdot d \ln \ell_g$ and reduces this group's wages by $\theta_{gg} \cdot (1 - s_g^y)$. The diagonal terms in the propagation matrix thus specify the slope of the aggregate elasticity of demand for group g . The supply shift alters other groups' wages by $d \ln w_{g'} = (\frac{1}{\lambda} \cdot s_g^y - \theta_{g'g}) \cdot d \ln \ell_g$ and we can see that g and g' are q -complements if $\frac{1}{\lambda} > \frac{1}{s_g^y} \cdot \theta_{g'g}$

(or equivalently, from symmetry $\frac{1}{\lambda} > \frac{1}{s_{g'}^y} \cdot \theta_{gg'}$). Pairs of groups with large corresponding off-diagonal entries can be q -substitutes. With the standard CES aggregate production function (with a common elasticity of substitution), all groups are q -complements.

6 The multi-sector economy

In this section, we generalize our results to a multi-sector economy. The multi-sector extension is important for several reasons. First, the way we measure direct task displacement in the rest of the paper relies on this extension since, in reality, the rate at which tasks are automated varies substantially across industries. Second, the multi-sector economy enables us to incorporate the consequences of a richer menu of competing technological effects—including those that work through industry-level productivity shocks—and the implications of changes in markups.

6.1 Environment

A (unique) final good y is produced by combining the output y_i of a finite number of industries, indexed by $i \in \mathbb{I} = \{1, 2, \dots, I\}$, via a constant returns to scale function $y = f(y_1, \dots, y_I)$. We denote the unit-cost function for the final good by $c^f(p)$, where $p = (p_1, \dots, p_I)$ is the vector of sector prices. We also denote the share of industry i in the economy by $s_i^y(p) = \partial \ln c^f(p) / \partial \ln p_i$, which depends on the vector of sector prices (where the equality is a consequence of Shephard's lemma). We continue to set the final goal as the numeraire.

Production in each sector y_i requires the completion of the tasks in the set \mathcal{T}_i , where \mathcal{T}_i has positive measures given by $M_i > 0$. We assume without loss of generality that the sets $\{\mathcal{T}_i\}_{i \in \mathbb{I}}$ are disjoint and denote their union by \mathcal{T} , which makes up the tasks space of the entire economy.²⁴ As in our one-sector setup, task quantities $y(x)$ are aggregated using a constant elasticity of substitution (CES) aggregator with elasticity $\lambda \in (0, 1)$:

$$y_i = A_i \cdot \left(\frac{1}{M_i} \int_{\mathcal{T}_i} (M_i \cdot y(x))^{\frac{\lambda-1}{\lambda}} dx \right)^{\frac{\lambda}{\lambda-1}},$$

where the new term, A_i , is a Hicks-neutral sector-specific productivity term.

An additional new element is that we allow for exogenous sector-specific markups, denoted by $\mu_i \geq 1$. This assumption allows us to model labor market implications of changing markups within the US economy (as studied, for example, in De Loecker et al., 2020). The case with $\mu_i = 1$ for all $i \in \mathbb{I}$ is a special case corresponding to a competitive economy.

As in the one-sector model, tasks are produced according to (1). We continue to assume that labor is inelastically supplied while the capital needed for any task $x \in \mathcal{T}$ is produced from the final good at a constant marginal cost of 1.

We also continue to impose Assumption 1 from the one-sector model, except that the finite integrals and strict comparative advantage are now imposed sector by sector.

²⁴ It is straightforward to allow for the same tasks to be performed in different industries, and whether we do so or not has no relevance for the results below.

6.2 Equilibrium

A market equilibrium is given by a positive vector of real wages $w = \{w_g\}_{g \in \mathbb{G}}$, a positive vector of sectoral prices $p = \{p_i\}_{i \in \mathbb{I}}$, an aggregate output level y , an allocation of tasks to skill groups $\{\mathcal{T}_{gi}\}_{g \in \mathbb{G}, i \in \mathbb{I}}$ and capital $\{\mathcal{T}_{ki}\}_{i \in \mathbb{I}}$ in each industry, task prices $\{p(x)\}_{x \in \mathcal{T}}$, task labor demands $\{\ell_g(x)\}_{g \in \mathbb{G}, x \in \mathcal{T}}$ and capital production levels $\{k(x)\}_{x \in \mathcal{T}}$ such that:

E1 Task prices are equal to the minimum unit cost of producing the task:

$$p(x) = \min \left\{ \frac{1}{A_k \psi_k(x)}, \left\{ \frac{w_g}{A_g \psi_g(x)} \right\}_{g \in \mathbb{G}} \right\}.$$

E2 Tasks are produced in a cost-minimizing way, which means that for each sector $i \in \mathbb{I}$, the set of tasks

$$\mathcal{T}_{gi}(w) = \left\{ x : p(x) = \frac{w_g}{A_g \psi_g(x)} \right\}$$

is allocated to workers from skill group $g \in \mathbb{G}$, and the set of tasks

$$\mathcal{T}_{ki}(w) = \left\{ x : p(x) = \frac{1}{A_k \psi_k(x)} \right\}$$

is produced with capital (where we condition on the vector of wages for later reference).

E3 Task-level demands for labor (for any $g \in \mathbb{G}$) and capital are given by

$$\ell_g(x) = \begin{cases} y_i \cdot p_i^\lambda \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \frac{1}{M_i} \cdot A_g^{\lambda-1} \cdot \psi_g(x)^{\lambda-1} \cdot w_g^{-\lambda} & \text{for } x \in \mathcal{T}_{gi}(w) \\ 0 & \text{otherwise.} \end{cases}$$

and

$$k(x) = \begin{cases} y_i \cdot p_i^\lambda \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \frac{1}{M_i} \cdot A_k^{\lambda-1} \cdot \psi_k(x)^{\lambda-1} & \text{for } x \in \mathcal{T}_{ki}(w) \\ 0 & \text{otherwise.} \end{cases}$$

E4 The labor market clears for all g :

$$\sum_i \int_{\mathcal{T}_{gi}} \ell_g(x) \cdot dx = \ell_g.$$

E5 Sector i 's price is given by its marginal cost times markup μ_i :

$$p_i = \mu_i \cdot \frac{1}{A_i} \cdot \left(\frac{1}{M_i} \int_{\mathcal{T}_i} p(x)^{1-\lambda} \cdot dx \right)^{1/(1-\lambda)}.$$

E6 The price of the final good is 1, which implies

$$1 = c^f(p).$$

In addition, as in the one-sector model, we use the tie-breaking rule that when a task can be performed at equal cost by multiple factors, it is first assigned to capital and then to lower-indexed skill groups ahead of higher-indexed groups. Strict comparative advantage again ensures that such ties can occur only on a set of measures zero, and thus this tie-breaking rule is inconsequential.

Fig. 5 provides a graphical illustration of the equilibrium, emphasizing the allocation of the tasks in each industry to different factors and their aggregation to the production of the unique final good.

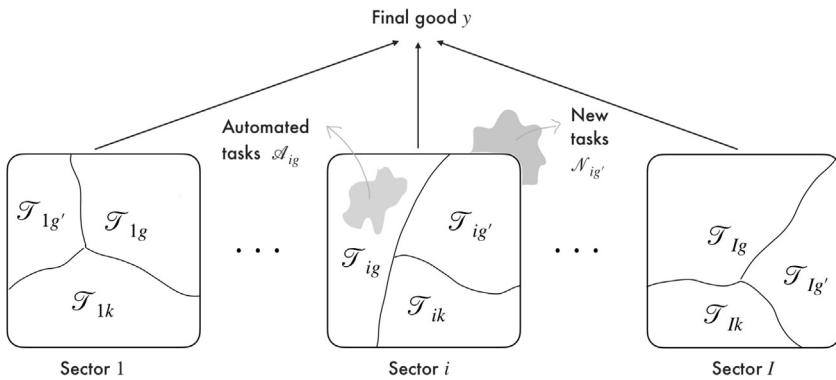


FIGURE 5 Equilibrium task assignment and task shares. The figure depicts the task space of a multi-sector economy and shows automation and new tasks taking place in industry i .

Most of these equilibrium conditions are familiar from the one-sector model. E1-E2 are identical to before and leverage cost-minimization. E3 and E5 are different from before because of the presence of markups: the latter condition imposes that industry prices incorporate markups, and the former adjusts factor demands for the presence of markups—higher markups translate into lower factor demands. E4 aggregates the demand for labor across industries, while E6 is again the numeraire condition.

As before, we can represent the equilibrium in terms of task shares, but now defined separately by sector $i \in \mathbb{I}$:

$$\begin{aligned}\Gamma_{gi}(w) &\equiv \frac{1}{M_i} \int_{\mathcal{T}_{gi}(w)} \psi_g(x)^{\lambda-1} \cdot dx \text{ for } i \in \mathbb{I} \text{ and } g \in \mathbb{G} \\ \Gamma_{ki}(w) &\equiv \frac{1}{M_i} \int_{\mathcal{T}_{ki}(w)} \psi_k(x)^{\lambda-1} \cdot dx \text{ for } i \in \mathbb{I}.\end{aligned}$$

Proposition 8 (Equilibrium representation). *Equilibrium wages w , industry prices p , and level of output y , solve the system of equations*

$$w_g = \left(\frac{y}{\ell_g} \right)^{1/\lambda} \cdot A_g^{1-1/\lambda} \cdot \left[\sum_i \bar{s}_i^y(p) \cdot p_i^{\lambda-1} \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \Gamma_{gi}(w) \right]^{1/\lambda} \quad \text{for } g \in \mathbb{G}, \quad (25)$$

$$p_i = \mu_i \cdot \frac{1}{A_i} \cdot \underbrace{\left(\Gamma_{ki}(w) \cdot A_k^{\lambda-1} + \sum_g \Gamma_{gi}(w) \cdot \left(\frac{w_g}{A_g} \right)^{1-\lambda} \right)^{1/(1-\lambda)}}_{\equiv \mathcal{C}_i(w)} \quad \text{for } i \in \mathbb{I}, \quad (26)$$

$$1 = c_f(p), \quad (27)$$

where $\mathcal{C}_i(w)$ denotes the marginal cost of producing output of sector i .

This characterization is analogous to the one in Proposition 1 for the one-sector model, except that we now also have an additional equilibrium condition for sectoral prices.

6.3 Effects of technology in the multi-sector economy

We can use the characterization in Proposition 8 to derive the effects of different types of technologies on the equilibrium wage structure. To do this, we rely again on the propagation matrix, which in this case can be written as

$$\Theta = \frac{1}{\lambda} \cdot \left(1 - \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \right)^{-1},$$

where the Jacobian $\partial \ln \Gamma(w) / \partial \ln w$ is now the $G \times G$ matrix with its gg' th entry given by

$$\sum_i \omega_{gi} \cdot \frac{\partial \ln \Gamma_{gi}(w)}{\partial \ln w_{g'}},$$

where ω_{gi} denotes wage payments received by group g in industry i as a share of total group wage payments. This matrix summarizes how changes in the wage of group g' affects group g by summing over the effects taking place in different industries.

As in the previous section, we start with the direct effects of new technologies, represented by the vector z , on the demand for skill group g . Define z_{gi} as the percent change in demand for workers from group g in industry i due to a change in technology at constant factor and sectoral prices. For automation, new tasks, and augmenting technologies, this coincides with the effects of these technologies on workers' task shares in industry i . We also define the

productivity gains at the sectoral level from these technological advances as $\pi_i = -d \ln \mathcal{C}_i(w)|_{w=\text{constant}} > 0$, which summarizes the contribution of technology to TFP in sector i . Finally, to simplify the exposition, we assume industries are combined into the final good with a constant elasticity of substitution η , though this can be relaxed.

Proposition 9 (Effects of technology in the multi-sector economy). *Consider a change in technology with direct effects $\{z_{gi}\}_{g \in \mathbb{G}, i \in \mathbb{I}}$ and productivity gains $\{\pi_i\}_{i \in \mathbb{I}}$. The effect of this technology on wages, sectoral prices, and output is given by*

$$d \ln w = \Theta \cdot \text{stack} \left(d \ln y + \sum_i \omega_{gi} \cdot z_{gi} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right) \quad (28)$$

$$d \ln p_i = \sum_g s_g^{y_i} \cdot d \ln w_g - \pi_i \quad \text{for } i \in \mathbb{I} \quad (29)$$

$$0 = \sum_i s_i \cdot d \ln p_i. \quad (30)$$

Here $s_g^{y_i}$ is the share of payments to skill group g in value-added in industry i , s_i is the share of industry i in total costs, and ω_{gi} denotes wage payments received by group g in industry i as a share of total group wages.

The proposition decomposes the effects of technology on the wage structure into four distinct channels. The first is the productivity effect, represented by $d \ln y$. The second comprises the usual direct effects of technology, the z_{gi} 's, except that these are now at the industry level and have to be aggregated. The third is captured by the propagation matrix, Θ , pre-multiplying the vector on the right-hand side of Eq. (28), which again summarizes the role of ripple effects.

The fourth and new element is the last term on the right side of (28). This corresponds to changes in the sectoral composition of the economy, which can be non-neutral if expanding sectors differ from contracting ones in their factor demands. Conversely, these changes are neutral when all sectors employ the same input mix. More generally, this term captures two forces. On the one hand, a reduction in the price of sector i increases its quantity, raising its demand for labor. This sectoral-demand effect depends on the elasticity of substitution between sectors (assumed to be equal to η). On the other hand, a reduction in the price of sector i reduces the value of the marginal product of workers and the demand for their services with an elasticity λ . When $\lambda > \eta$, the first effect dominates, and sectoral shifts benefit workers in sectors experiencing less productivity growth. This captures the same economic mechanism as in the celebrated Baumol effect (Baumol et al., 2012): workers specializing in sectors with lower (technological) productivity growth, such as healthcare, tend to benefit because the relative prices of these sectors increase strongly as aggregate output expands.

Finally, the exact equilibrium changes in sectoral prices can be obtained from (29), while Eq. (30) pins down the change in the output level.

It is useful to illustrate the results of Proposition 9 for automation technologies. The effects of automation on wages are now given by

$$d \ln w = \Theta \cdot \text{stack} \left(d \ln y - d \ln \Gamma_g^{\text{auto}} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right), \quad (31)$$

with $d \ln \Gamma_g^{\text{auto}}$ the total direct task displacement due to automation experienced by group g ,

$$d \ln \Gamma_g^{\text{auto}} = \sum_i \omega_{gi} \cdot d \ln \Gamma_{gi}^{\text{auto}}. \quad (32)$$

This is obtained by summing the direct task displacement from automation experienced by group g in industry i , $d \ln \Gamma_{gi}^{\text{auto}}$, across industries. The summation weights are given by the shares of wage payments from industry i in group g 's total wage payments. The wage equation (31) again contains the usual productivity and displacement effects of automation, as well as the ripples via the propagation matrix.

The new element here relative to the single-sector economy is the indirect effect of automation working via its impact on sectoral prices, which shifts the composition of the economy. These effects depend on the contribution of automation to the TFP of the different sectors, which is given by $\pi_i = \sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{auto}} \cdot \pi_{gi}^{\text{auto}}$, where the π_{gi}^{auto} 's are the average cost savings from automation in sector i . For $\lambda > \eta$, which is the case we consider in our quantitative exercise, automation reallocates labor demand away from sectors that automate at a higher rate, reducing the relative wages of workers in these industries.

The equilibrium here is not generically efficient because of the presence of markups. Nevertheless, when there are no markups or when markups are uniform across sectors ($\mu_i = \mu$), the equilibrium is again efficient. In that case, Eqs. (29) and (30) imply that average wage changes from automation are

$$\sum_g s_g^y \cdot d \ln w_g = \underbrace{\sum_i s_i \sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{auto}} \cdot \pi_{gi}^{\text{auto}}}_{=d \ln tfp},$$

where the term on the right-hand side is aggregate TFP, obtained by summing the cost savings due to automation in different industries. As in the single-sector model, we can see the effect of automation on wage levels depends on its contribution to TFP, and could be large or small depending on how big the cost savings due to this technology are.

6.4 Sectoral TFP and markups

The multi-sector economy allows us to study the labor market implications of sector-specific (Hicks-neutral) technological advances and changes in markups. In particular, Proposition 9 also applies to sector-specific technologies, which are important drivers of structural change in the economy (see Ngai and Pissarides, 2007; Buera et al., 2021). The effect of these technologies satisfies

$$\begin{aligned} d \ln w &= \Theta \cdot \text{stack} \left(d \ln y - (1 - \lambda) \cdot \sum_i \omega_{gi} \cdot d \ln A_i + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right) \\ d \ln p_i &= \sum_g s_g^{y_i} \cdot d \ln w_g - d \ln A_i \quad \text{for } i \in \mathbb{I} \\ 0 &= \sum_i s_i \cdot d \ln p_i. \end{aligned}$$

Hicks-neutral increases in sectoral TFP affect the wage structure via the four channels identified above. The first is the productivity effect, which corresponds to the expansion of output, $d \ln y$. The second works through the reduction in task prices for the sectors that become more productive. Task-price effects are aggregated according to the exposure of different skill groups to the industry in question, as measured by the wage-bill shares ω_{gi} . The third channel is via the ripple effects, encoded in the propagation matrix Θ . The fourth is the sectoral price changes in the last term on the right-hand side of the wage equation.

The comparison of this wage equation to (31) shows the differences between sectoral TFP improvements and automation. Automation works via the extensive-margin of task reallocation taking place within sectors, while there is no equivalent of these effects in the case of sectoral TFP, which works by reallocating labor demand across sectors.

Finally, we can derive the effects of changes in markups. This follows from our characterization of the equilibrium in Proposition 8 and is presented next.

Proposition 10 (Effects of markups in the multi-sector economy). *Consider an exogenous change in sectoral markups $\{d \ln \mu_i\}_{i \in \mathbb{I}}$. The impact on wages, sectoral prices, and output is given by*

$$d \ln w = \Theta \cdot \text{stack} \left(d \ln y - \lambda \cdot \sum_i \omega_{gi} \cdot d \ln \mu_i + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right) \quad (33)$$

$$d \ln p_i = \sum_g s_g^{y_i} \cdot d \ln w_g + d \ln \mu_i \quad \text{for } i \in \mathbb{I} \quad (34)$$

$$0 = \sum_i s_i \cdot d \ln p_i. \quad (35)$$

This proposition shows that markups affect the wage structure via the same four channels identified in Proposition 9. The first is the productivity effect, given by $d \ln y$, which results from the fact that increases in markups can reduce output. The second is the direct impact of the changes in markups, which are aggregated using wage-bill shares. This effect is negative because markups reduce (relative) production in the affected sectors. The third is through the ripple effects that these changes induce, which work via the propagation matrix, Θ , as characterized above. The fourth channel is the shifts in the sectoral composition of the economy due to markups.

Just like the sector-specific technology terms discussed above, markups' impact all workers in an industry uniformly. This is why their distributional effects work through shifts in labor demand across sectors—and they do not generate any displacement or reinstatement. The distributional effects of this reallocation across sectors will be muted if expanding and contracting sectors do not differ substantially in their skill mixes. This is the reason why we expect, from a theoretical point of view, these effects to be less pronounced than those coming from automation and new tasks, and this is indeed what we document in our empirical application, presented next.

7 Reduced-form evidence

In this section, we estimate reduced-form equations derived from the task framework. We focus on US labor markets between 1980 and 2016. The estimates support the key prediction of the task framework, showing that *extensive-margin* changes in the allocation of tasks to factors, driven by automation and new tasks, have first-order effects on the wage structure. In fact, these effects are much larger than those estimated for other technologies. Consistent with the expectation that automation and new tasks shift labor demand, we find that these forces have had large impacts on employment outcomes as well.

We first summarize the trends in wages and employment that we seek to explain. We then derive our reduced-form specification and discuss how we measure the displacement due to automation and reinstatement due to new tasks experienced by US worker groups.

7.1 US labor market trends

Fig. 6 depicts the major wage inequality trends in the US. It plots cumulative real hourly wage growth since 1960 by gender (separately in the left and the right panels) and education level. We show data from the CPS (with connected dots) and the decennial Censuses and the ACS (with diamonds). In the 1960s and 70s, hourly wages grew by 1.5%-2% per year for all groups, and the real wage growth tracked labor productivity, implying that the labor share of national income remained stable. From 1980 to 2016, we see a strikingly different pattern: hourly real wages continue to grow for workers with a college degree and even more so for those with a postgraduate degree, while wages for noncollege

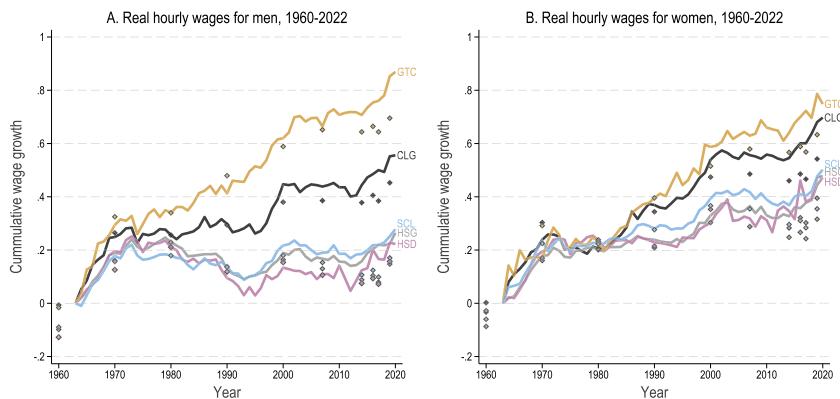


FIGURE 6 Cumulative growth in real hourly wages for men and women by education level (GTC: postgraduate degree, CLG: college degree, SCL: some college, HSG: high school degree, HSD: high school dropout), 1960–2022. Diamonds: data from the US Census and the American Community Survey. Connected line: data from the Current Population Survey. Wages deflated using the personal consumption expenditure index from the Bureau of Economic Analysis.

workers stagnated and, for men with a high school degree or less, even declined in some periods.

In line with the tepid wage growth observed during this period, the labor share of national income declined markedly since 1980, as shown in Fig. 7, especially in manufacturing and retail.

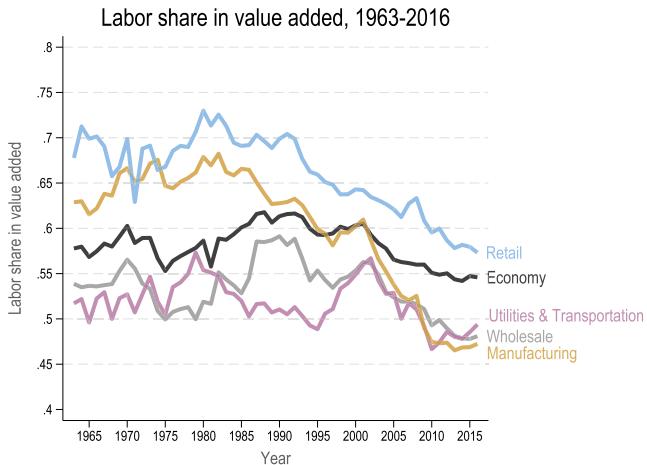


FIGURE 7 The evolution of labor shares in manufacturing, wholesale, retail, utilities and transportation. Data from the BEA-BLS Integrated Industry Accounts, 1963–2016.

These unequal wage trends coincided with rising disparities in employment rates, shown in Fig. 8. Since 1980, employment rates for college-graduate men have remained stable, and employment rates for college graduate women have

continued to increase. At the same time, employment rates for men without a college degree declined (though the beginning of this trend dates to the 1970s), while employment rates among women with a high school degree and an associate degree decelerated and started to decline.

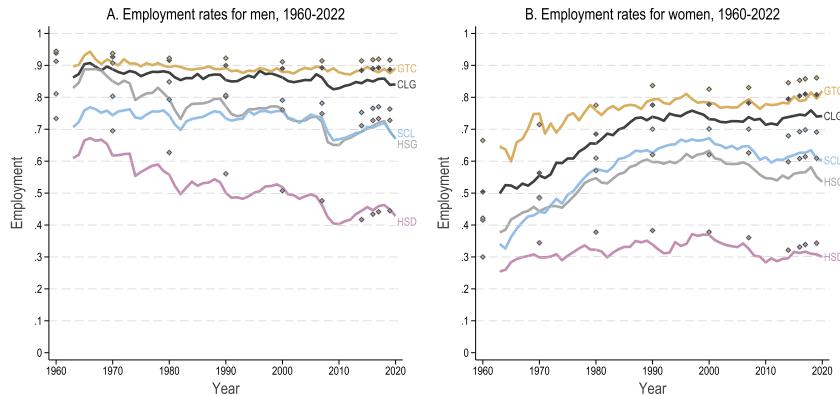


FIGURE 8 Employment rates for men and women by education level (GTC: postgraduate degree, CLG: college degree, SCL: some college, HSG: high school degree, HSD: less than high school), 1960–2022. Data from the US Census and the American Community Survey are shown as diamonds, and data from the Current Population Survey are shown as the connected lines.

7.2 Specification

For our reduced-form analysis, we organize the data at a more granular level than in Figs. 6 and 8 and look at 500 demographic groups, proxying for skill groups in our theory. These demographic groups are defined by the five education levels, gender, five age groups (16–25 years of age, 26–35, 36–45, 46–55, 56–65), ethnicity (White, Black, Hispanic, and Asian), and native vs. foreign-born status. For each group, we compute the change in log hourly wages and the change in log hours worked from 1980 to 2016 using the 1980 Census and pooling five years of the American Community Survey (ACS) between 2014 and 2018. Our reduced-form specification relates wage changes experienced by groups between 1980 and 2016 to proxies of automation, new tasks, sectoral TFP growth and markups.

To motivate our specification, start from Eq. (28) and rewrite it as

$$d \ln w_g = \theta_{gg} \cdot \left(d \ln y + \sum_i \omega_{gi} \cdot z_{gi} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right) + \text{Ripple effects}_g, \quad (36)$$

where the Ripple effects_g term captures spillovers from shocks impacting other worker groups. Our reduced-form analysis treats the ripple effects as part of the

error term and focuses exclusively on the relationship between shocks directly affecting a group and its outcomes. In addition, we assume that the diagonal entries $\theta_{gg} = \theta$ are equal, which motivates a simple linear model for group wage changes. We estimate the ripple effects in Section 8.

Our specification accounts for various technologies affecting labor demand directly, via the term $\sum_i \omega_{gi} \cdot z_{gi}$ above. First, we consider the role of automation, whose direct effect $d \ln \Gamma_g^{\text{auto}}$, was defined in (32) as the summation of industry-level task displacements $d \ln \Gamma_{gi}^{\text{auto}}$'s across industries. Second, we consider the role of new tasks, whose direct effect is to reinstate group g . This reinstatement effect is also given by a summation across industries:

$$d \ln \Gamma_g^{\text{new}} = \sum_i \omega_{gi} \cdot d \ln \Gamma_{gi}^{\text{new}}.$$

Additionally, we assume that labor-augmenting technologies, $d \ln A_g$, satisfy

$$d \ln A_g = \delta_{\text{education}_g} + \delta_{\text{gender}_g} + u_g,$$

where u_g is a residual independent of other covariates. The term $\delta_{\text{education}_g}$ incorporates common improvements in labor productivity that apply to all workers with the same education level. This formulation is similar to but more general than those typically considered in the skill-biased technical change literature.²⁵ The term δ_{gender_g} allows for shifts in technology or labor market discrimination affecting women relative to men.

The resulting estimation equation is

$$\begin{aligned} \Delta \ln w_g = & \text{constant} + \beta^{\text{auto}} \cdot \text{Task displacement from automation}_g^{1980-2016} \\ & + \beta^{\text{new}} \cdot \text{Task reinstatement from new tasks}_g^{1980-2016} \\ & + \text{Dummies for education level} + \text{Dummies for gender} \\ & + \beta^{\text{sector}} \cdot \text{Sectoral shifts}_g + \underbrace{\text{Ripples}_g + u_g}_{=v_g}, \end{aligned} \quad (37)$$

where we rewrote Eq. (36) for wage changes between 1980 and 2016. In this regression model, the productivity effect, $d \ln y$, is included in the constant. We also replaced $d \ln \Gamma_g^{\text{auto}}$ and $d \ln \Gamma_g^{\text{new}}$ with their empirical counterparts, whose construction we discuss below. The education and gender dummies are included to account for the common shifts in labor-augmenting technology for all workers of a given education level or gender, as explained above. The error term is then

²⁵ One could also introduce changes in labor-augmenting technology at the intensive margin in a similar way. As with uniform changes, any increase in labor-augmenting technology at the intensive margin (the $d \ln \psi_g^{\text{intensive}}$'s) that is common across educational groups would be subsumed by education fixed effects, and residual changes would be part of the error term in the estimation equation.

a combination of the ripple effects and residual changes in group-level productivity. We present estimates that condition on education and gender dummies and estimates that do not, which allows us to explore the extent to which the reduced-form model can explain the observed wage trends between and within educational groups and gender.

Our regression model also includes extra terms to control for sectoral shocks and changes in sectoral composition. As a first strategy, in some of our regression models, we account for the influence of changes in sectoral markups and TFP on the wage structure. Building on the analysis in Section 6, the influence of these forces on the wage structure can be expressed as

$$\text{Sectoral TFP}_g = \sum_i \omega_{gi} \cdot \Delta \ln \text{Multifactor TFP}_i,$$

$$\text{Sectoral markups}_g = \sum_i \omega_{gi} \cdot \Delta \ln \text{Markups}_i,$$

where ω_{gi} is the share of wages group g receives from industry i (computed using the 1980 Census), $\Delta \ln \text{Multifactor TFP}_i$ is the change in industry TFP over 1980–2016 (computed for 50 industries using the BEA-BLS Integrated Industry Accounts, which are then matched to the 1980 Census), and $\Delta \ln \text{Markups}_i$ are estimates of the average markup change for these industries (taken from Hubmer and Restrepo, 2021).

As a second strategy, we follow Acemoglu and Restrepo (2022) and explore regression models that directly control for the observed changes in sectoral value added, including a term of the form

$$\text{Sectoral value-added shares}_g = \sum_i \omega_{gi} \cdot \Delta \ln \text{Value-added share}_i,$$

where $\Delta \ln \text{Value-added share}_i$ is the change in industry value added over 1980–2016 (computed from the BEA-BLS Integrated Industry Accounts). This control captures the influence of *all* observed shifts in the sectoral composition of the US economy during this period, including changes induced by automation and new tasks, on group g 's wages. For this reason, these estimates remove any indirect effects of automation and new tasks working through changes in the sectoral composition of the economy (i.e., the term $(\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i$ in Proposition 9). Our quantitative exercise in Section 8 returns to this issue and shows that these indirect effects of automation and new tasks are estimated to be small during this period.

Besides the regression model in (37), we estimate equations with changes in log hours worked per person in each group as outcome. Since the technology terms on the right side of (37) shift labor demand, we expect them to impact employment in the same direction.

7.3 Measuring automation and new tasks

As in Acemoglu and Restrepo (2022), we measure task displacement due to automation using automation-induced industry labor share changes and information on which types of workers within an industry are likely to be impacted by automation. We assume that automation in an industry only displaces workers in routine occupations and that such displacement takes place at equal rates for workers in these occupations, regardless of their groups. This means that if there are workers from two demographic groups g and g' in a routine occupation undergoing automation, then the same proportion of workers from these two demographic groups in this occupation will be displaced.

Under these assumptions, we show in the Appendix that task displacement due to automation in industry i can be obtained as²⁶:

$$d \ln \Gamma_{gi}^{\text{auto}} = \text{RCA routine}_{gi} \cdot (-\Delta \ln s_L^{y_i, \text{auto}}). \quad (38)$$

Combining this with (32), our measure of total task displacement experienced by group g is:

$$\begin{aligned} & \text{Task displacement from automation}_g^{1980-2016} \\ &= \sum_i \omega_{gi} \cdot \text{RCA routine}_{gi} \cdot (-\Delta \ln s_L^{y_i, \text{auto}}), \end{aligned} \quad (39)$$

where

- RCA routine_{gi} is the revealed comparative advantage of group g in routine tasks in industry i . This term adjusts for the incidence of automation across workers in an industry. Intuitively, if group g performs all routine tasks in industry i , then an increase in automation in that industry will displace group g only. If multiple groups perform routine tasks in the industry, then an increase in automation in that industry will displace them in proportion to the share of routine tasks they perform in that industry (which our revealed comparative advantage captures). This term is computed from the 1980 Census as the ratio of wages earned by group g in routine jobs in industry i divided by all wage payments in routine jobs in the industry. We define routine jobs as the top one-third of occupations with the highest routine content, using the measure of routine tasks from ONET from Acemoglu and Autor (2011).
- $-\Delta \ln s_L^{y_i, \text{auto}}$ is automation-induced percent reduction in labor share in industry i . This term corresponds to the total share of tasks lost to automation among all workers in the industry. This automation-induced change in labor share is computed in two steps. In the first step, we run a regression of the observed percent decline in industry labor shares from 1987 to 2016 from the BEA-BLS integrated industry accounts against three proxies of automation.

²⁶ This formula is exact for $\lambda = 1$. The general case with non-unitary elasticity of substitution between tasks includes an additional adjustment term, but does not appreciably change the results, as we further discuss in the next section.

These proxies include the adjusted penetration of robots over 1993–2007, computed from European countries that are ahead of the US in terms of robot adoption and incorporating the adjustment discussed in Acemoglu and Restrepo (2020a); the change in expenditure on dedicated machinery divided by industry value added, 1987–2016, and the change in expenditure on specialized software divided by industry value added, 1987–2016 (the latter two sourced from the BLS detailed capital tables). These regressions are reported in Acemoglu and Restrepo (2022) and show that these three proxies account for 50% of the cross-industry variation in labor shares. In a second step, we take the predicted labor share change from this cross-industry regression and use it as a measure of the labor share decline driven by automation.

Fig. 9 summarizes the results of this measurement exercise. It depicts both the observed labor share declines and the predicted declines driven by automation (both in percent terms, and the former in blue and the latter in orange). Observed labor share declines and those driven by automation are highly correlated, but there are also some notable exceptions. Several industries that are part of the transport sector have large overall declines in labor share, but only moderate predicted declines due to automation—because they have relatively low levels of robot penetration and small changes in dedicated machinery and specialized software expenditures. Several other industries, including automobile manufacturing, show both sizable observed declines and predicted declines.²⁷

- ω_{gi} is group g 's exposure to industry i , which is used as weight in summing across industry-level task displacements. This term captures the importance of tasks performed in industry i for group g , and is computed from the 1980 Census for 50 industries that we can track consistently in the BEA-BLS integrated industry accounts.

Our measure of reinstatement due to new tasks uses data from Lin (2011), which are also analyzed in Acemoglu and Restrepo (2018b). These data, in turn, rely on new job titles from the Dictionary of Occupational Titles (DOT) in 1977 and 1991 and from the 2000 Census. Using these data, we construct task reinstatement for group g in industry i as

$$\begin{aligned} d \ln \Gamma_{gi}^{\text{new}} = & \sum_o \omega_{gio}^{1980} \cdot \text{Share new job titles DOT 1977} \\ & + \sum_o \omega_{gio}^{1990} \cdot \text{Share new job titles DOT 1991} \\ & + \sum_o \omega_{gio}^{2000} \cdot \text{Share new job titles Census 2000}, \end{aligned}$$

²⁷ One could use these proxies directly as regressors or instruments, and we do this in Acemoglu and Restrepo (2022). Projecting these measures on the labor share decline is helpful because it converts them into units of “tasks lost” to automation and allows us to summarize their effects in a single variable representing the task displacement associated with these technologies.

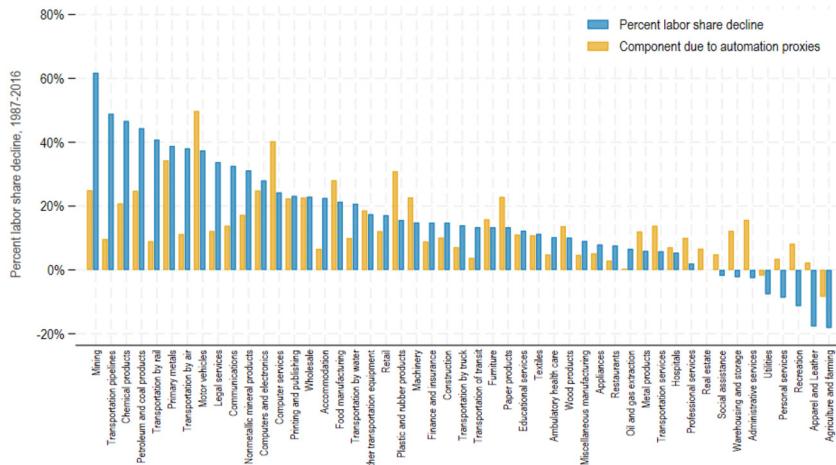


FIGURE 9 Percent decline in industry labor shares (in blue) and the predicted labor share declines due to automation (in orange). The observed declines are computed from the BEA-BLS Integrated Industry Accounts. The predicted declines are from a regression of the observed declines on the adjusted penetration of robots (from Acemoglu and Restrepo, 2020a), as well as the increases in expenditures in dedicated machinery and specialized software as a share of value added (both from the BLS Detailed Capital Tables).

where ω_{gio} denotes the share of total wage payments to group g in industry i that come from occupation o . Analogous to the total task displacement measure, total task reinstatement for group g is computed as

$$\text{Task reinstatement from new tasks}_g^{1980-2016} = \sum_i \omega_{gi} \cdot d \ln \Gamma_{gi}^{\text{new}}. \quad (40)$$

The assumption behind these measures is that new job titles proxy for new tasks (and are not just a relabeling of existing jobs), that each new job title has the same positive impact on new tasks, and that new tasks are proportionately spread among workers in the occupations in which they emerge. These considerations also motivate the use of the wage-bill share of different demographic groups in the occupation in the concurrent period to capture the importance of these new tasks for each group. We compute this measure using data for 300 detailed occupations that we can trace consistently over time and across Censuses and different waves of the ACS.²⁸

Before describing group-level patterns, we show in Fig. 10 that, at the occupational level, there is a strong positive association between new tasks (summed

²⁸ Notice that this is different from the measurement strategy of our baseline automation measure, which uses beginning of sample (1980) weights. This difference stems from the fact that, in the theory, new tasks benefit workers who end up taking over these tasks, while automation affects workers who used to work in the now-automated tasks. Tables A1 and A2 in the Appendix show that our reduced-form results are robust if we compute the new task measures using occupational shares fixed in 1980.

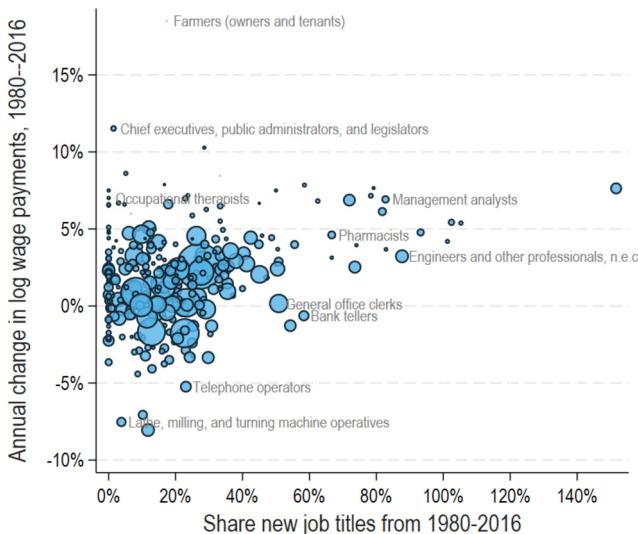


FIGURE 10 Changes in log total wage bill across occupations, 1980–2016 against share of new job titles introduced in each occupation (from DOT 1977, DOT 1991, and Census 2000). Data for 300 occupations.

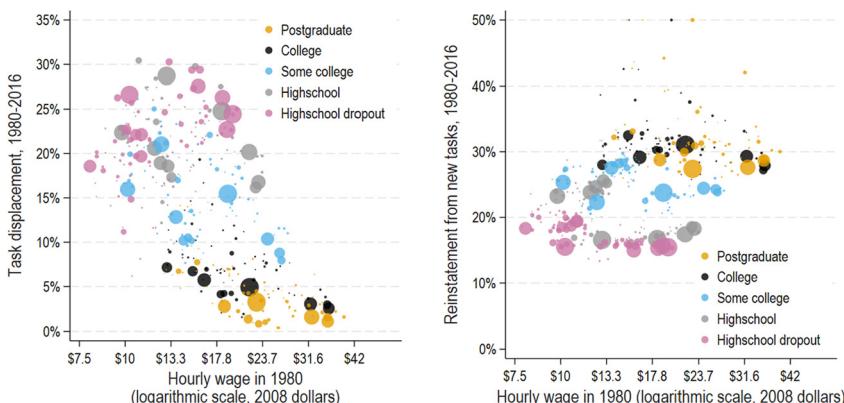


FIGURE 11 Left panel shows direct task displacement due to automation, 1980–2016 for 500 groups of US workers, and the right panel shows task reinstatement due to new job titles for these groups. Both panels plot these data against group average hourly wages in 1980, from the Census. Marker sizes are proportional to hours worked in 1980. Marker colors distinguish groups with different education levels.

over 1977, 1991 and 2000 measures) and labor demand. A 10 pp increase in job titles over this time window is associated with a 0.4 pp higher yearly growth rate of wage payments in that occupation from 1980 to 2016. This reproduces and extends the results reported in Acemoglu and Restrepo (2018b).

Fig. 11 provides a first comparison of our measures of task displacement from automation and reinstatement due to new tasks. The figure plots both vari-

ables against group-level hourly wages in 1980, to indicate where in the wage distribution the effects of displacement and reinstatement are felt. The left panel shows that, on average, US workers experienced a reduction in task shares of 19% during this period, but this was quite unevenly distributed in the population. While noncollege workers saw task share declines in the range of 20–30%, college and postgraduate workers were mostly shielded from such displacement.²⁹ The right panel, on the other hand, indicates that, on average, US workers benefited from a 22% expansion in their task shares due to new tasks. In contrast to automation, reinstatement effects are higher for more educated workers.

7.4 Reduced-form estimates

We begin by exploring the relationship between automation and labor market outcomes graphically. The top two panels in Fig. 12 provide bivariate scatter plots of the change in group wages from 1980–2016 (top left panel) and log hours per person (top right panel) against our measure of task displacement due to automation for this period. The bottom panel provides residual scatter plots that partial out education and gender dummies and sectoral value-added shares. Overall, the figure shows a negative association between task displacement due to automation and wage and employment changes. The associations are stable regardless of whether we include covariates. The estimated effects are also sizable. In the bottom panel, a 10 pp increase in task displacement for a skill group is associated with a 14.5% decline in wages and a 18.3% decline in hours worked relative to other groups.

Fig. 13 presents the analogous specifications for new tasks—with the top panel depicting the bivariate relationships and the bottom panel partialing out covariates. It shows a positive association between reinstatement due to new tasks and changes in wage and employment. The estimated effects are also sizable. In the bottom panel, a 10 pp more reinstatement due to new job titles is associated with a 17.6% increase in wages and a 14.0% increase in hours worked relative to other groups.

Figs. 12 and 13 support the key implications of the task framework: task displacement from automation is associated with negative wage consequences for exposed workers relative to others, while reinstatement due to new tasks is associated with positive wage effects. These technologies also have commensurate effects on employment—groups experiencing more task displacement have (relatively) lower hours worked, while the pattern is the opposite for those benefiting from greater task reinstatement.

Table 1 provides estimates for the change in log hourly wages as the outcome. Column 1 in Panels A and B report estimates of the bivariate relationships shown in the top-left panels of Figs. 12 and 13. The regression coefficient for

²⁹ Because this measure is based on predicted labor share declines over 1987–2016, we re-scale it to a 37-year equivalent change that matches the length of time used for the dependent variables (1980–2016).

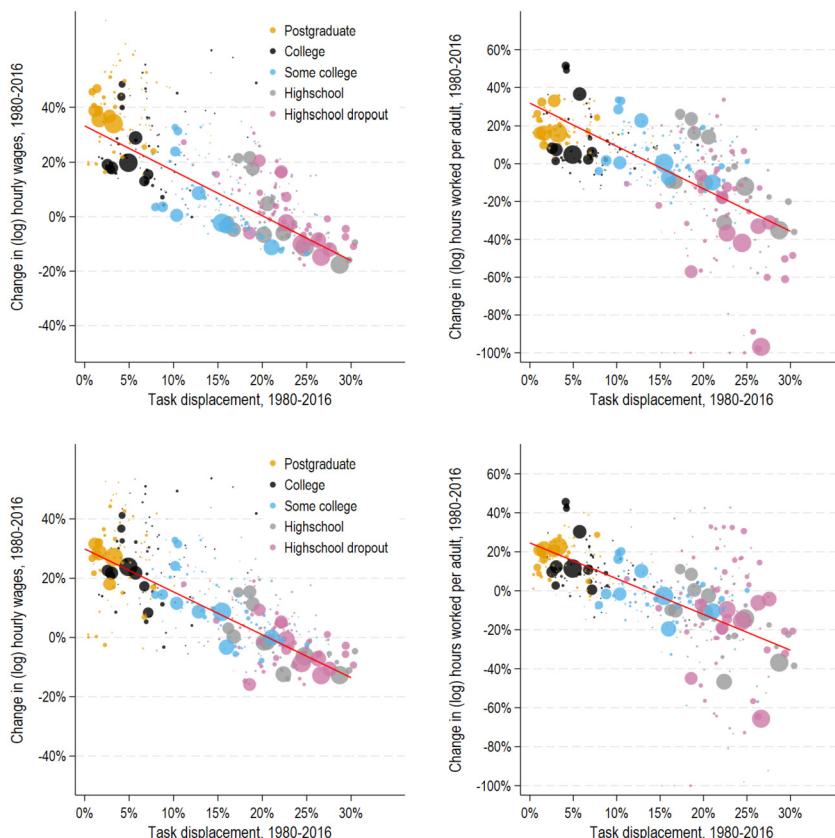


FIGURE 12 Reduced-form relationship between change in log hourly wages and change in log hours worked per person vs. task displacement due to automation, 1980–2016. The top panel presents bivariate scatter plots. The bottom panels present residual plots partialling out gender and education dummies and changes in sectoral value-added shares. Marker sizes are proportional to hours worked in 1980. Marker colors distinguish groups with different education levels.

task displacement is -1.65 (standard error = 0.10), while the coefficient for task reinstatement is 2.32 (standard error = 0.19).

Column 1 in Panel C includes both explanatory variables together. The coefficient for task displacement due to automation is now -1.19 (standard error = 0.23), and the coefficient for tasks reinstatement is 0.85 (standard error = 0.33). The point estimates are attenuated compared to Panels A and B, especially for new tasks, reflecting the fact that these two measures are negatively correlated. Nevertheless, these two variables jointly explain a remarkable 67% of the observed wage changes across worker groups in the US between 1980 and 2016, with automation accounting for 46% and new tasks for the remaining 20%.³⁰

³⁰ Throughout this section, we follow Klenow and Rodríguez-Clare (1997) and decompose the total R^2 into contributions from subsets of the variables by equally distributing the covariance terms

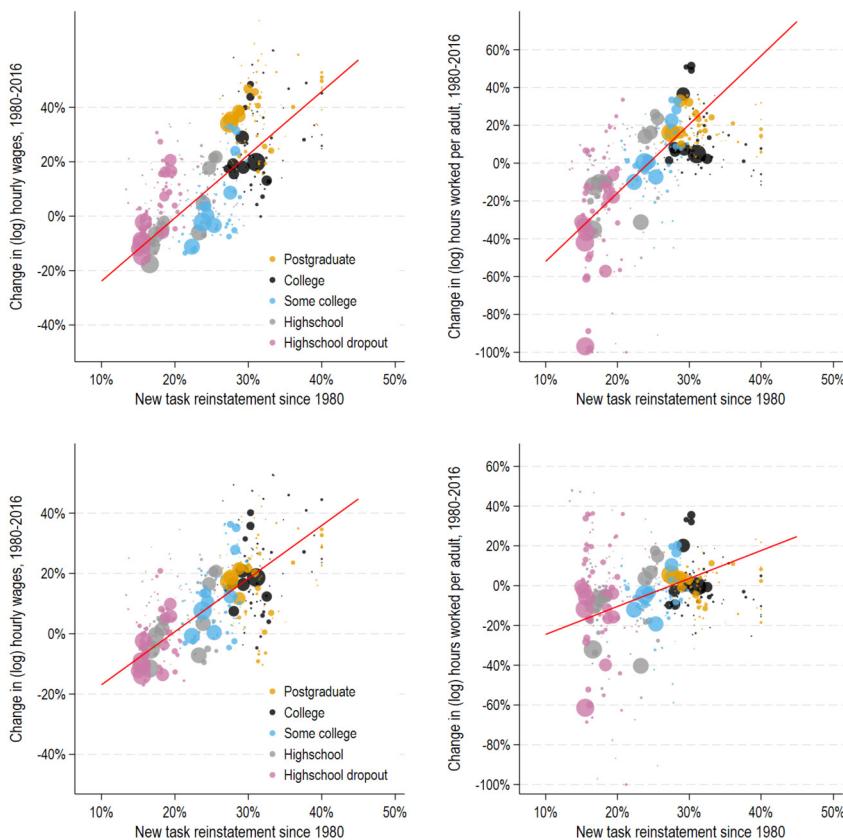


FIGURE 13 Reduced-form relationship between change in log hourly wages and change in log hours worked per person vs. reinstatement due to new tasks, 1980–2016. The top panel presents bivariate scatter plots. The bottom panels present residual plots partialling out gender and education dummies and changes in sectoral value-added shares. Marker sizes are proportional to hours worked in 1980. Marker colors distinguish groups with different education levels.

Note that these models do not include any other covariates, which means that our task displacement and reinstatement measures alone are responsible for the high explanatory power. Moreover, the high R^2 of these models shows that our task measures do an excellent job at accounting for the divergent wage trends across education and gender groups depicted in Fig. 6 and 8. This is because, as

between them. This means that the contribution of a covariate x_j to the explanatory power of a model of the form $y = \sum \beta_j x_j + u$ is

$$R^2 \text{ from } x_j = \beta_j \cdot \frac{\text{cov}(x_j, y)}{\text{var}(y)}.$$

By construction, these sum up to the model's total R^2 when added across all variables (subject to rounding).

highlighted in Fig. 10, our measures predict that non-college workers lost more tasks to automation and at the same time gain fewer new tasks than college educated workers.

The parameter estimates also imply sizable effects from both variables. A 10 pp increase in task displacement for a demographic is associated with 11.9% lower (relative) wages, while a 10 pp increase in task reinstatement is associated with 8.5% higher (relative) wages.

Panel D leverages the fact that task displacement and reinstatement are constructed to be in the same units and are thus predicted to impact wages with the same coefficient but opposite signs. This panel therefore combines these measures into a single explanatory variable, “net task change,” constructed as the difference between task reinstatement and displacement. This variable has a positive and precisely estimated coefficient, 1.05 (standard error = 0.07). Interestingly, this restriction only leads to a small reduction in the explanatory power of automation and new tasks, which, together, still account for 67% of the total variation in wage trends between demographic groups. This estimate implies that a 10 pp higher net task change is associated with a 10.5% increase in relative wages.

The remaining columns in Table 1 explore the robustness of these reduced-form relationships to the inclusion of various covariates. Column 2 controls for sectoral value-added shares, with little effect on the coefficient estimates for task displacement and reinstatement. Column 3 directly controls for sectoral shocks, and controls for changes in sectoral TFP and markups. The results are once more very similar, suggesting that automation and new tasks are distinct from these sectoral trends. More tellingly, we find that the sectoral variables explain 3–8% of the variance in wage trends across groups, while our task measures jointly explain 62–64%.

More importantly, Columns 4 and 5 add the education and gender dummies to the specifications from columns 2 and 3. In both specifications we continue to estimate a sizable negative association between group outcomes and automation and a substantial positive association with new tasks, with point estimates that are quite similar to those in column 1. Recall that education dummies account for the role of skill-biased (factor-augmenting) technologies benefiting more educated workers. These specifications thus suggest that automation and new tasks are distinct from these other forms of technological progress emphasized in previous literature. Moreover, the R^2 decomposition in these columns indicates that the explanatory power of education dummies is quite limited. The educational dummies (together with gender dummies and sectoral covariates) explain only 4–6% of the variance in wage trends across groups, while our task measures continue to jointly explain 78–80%. These decompositions imply that the extensive-margin changes associated with task displacement and reinstatement are more important drivers of wage trends between groups than the forces commonly emphasized in the literature and captured by the educational dummies and sectoral controls.

TABLE 1 Reduced-form evidence: changes in real hourly wages regressed on automation and new tasks, 1980–2016.

	Dependent variables: Change in log hourly wages, 1980–2016						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Only displacement from automation							
Automation task displacement	−1.65 (0.10)	−1.41 (0.20)	−1.50 (0.11)	−1.45 (0.18)	−1.41 (0.19)	−1.71 (0.25)	−1.75 (0.32)
R^2 for model	0.64	0.66	0.69	0.82	0.83	0.76	0.76
R^2 for automation	0.64	0.55	0.59	0.56	0.55	0.67	0.68
R^2 remaining covs		0.11	0.10	0.26	0.28	0.09	0.08
Observations	500	500	500	500	500	492	492
Panel B. Only reinstatement from new tasks							
New tasks reinstatement	2.32 (0.19)	2.09 (0.35)	2.37 (0.26)	1.76 (0.41)	1.56 (0.47)	2.18 (0.69)	2.94 (1.10)
R^2 for model	0.56	0.56	0.59	0.78	0.77	0.26	0.07
R^2 for new tasks	0.56	0.51	0.57	0.43	0.38	0.53	0.71
R^2 remaining covs		0.06	0.01	0.35	0.40	−0.27	−0.64
R^2 remaining covs	500	500	500	500	500	492	492
Panel C. Both explanatory variables							
Automation task displacement	−1.19 (0.23)	−1.18 (0.23)	−1.27 (0.22)	−1.28 (0.16)	−1.32 (0.17)	−1.55 (0.22)	−1.70 (0.29)
New tasks reinstatement	0.85 (0.33)	0.75 (0.38)	0.50 (0.37)	1.16 (0.32)	1.18 (0.37)	1.18 (0.36)	1.53 (0.47)
R^2 for model	0.67	0.67	0.69	0.84	0.84	0.77	0.76
R^2 for automation	0.46	0.46	0.50	0.50	0.51	0.60	0.66
R^2 for new tasks	0.20	0.18	0.12	0.28	0.29	0.28	0.37
R^2 remaining covs		0.03	0.08	0.06	0.04	−0.12	−0.27
Observations	500	500	500	500	500	492	492
Panel D. Net task change due to new tasks minus automation							
Net task change (new tasks–automation)	1.05 (0.07)	1.05 (0.14)	1.00 (0.08)	1.24 (0.13)	1.29 (0.17)	1.46 (0.19)	1.67 (0.29)
R^2 for model	0.67	0.67	0.69	0.84	0.84	0.76	0.75
R^2 for net task changes	0.67	0.66	0.63	0.78	0.81	0.92	1.06
R^2 remaining covs		0.00	0.05	0.06	0.03	−0.16	−0.30
R^2 remaining covs	500	500	500	500	500	492	492
<i>Other covariates:</i>							
Sectoral value added		✓		✓		✓	
Sectoral TFP			✓		✓		✓
Sectoral markups			✓		✓		✓
Gender and education dummies				✓	✓	✓	✓
Labor supply shifts						✓	✓

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hourly wages across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The dependent variable is the change in log hourly wages for each group between 1980 and 2016. Panel A reports results using only our task displacement measure. Panel B only uses our task reinstatement measure. Panel C includes both task displacement and task reinstatement on the right-hand side. Panel D combines task displacement and reinstatement into a net task change measure. The bottom rows list additional covariates included in each specification. In columns 6 and 7, we instrument changes in labor supply using changes in total hours worked by group from 1970 to 1980. All regressions are weighted by total hours worked by each group in 1980, as in Acemoglu and Restrepo (2022). Standard errors robust to heteroskedasticity are reported in parentheses.

Finally columns 6 and 7 control for labor supply changes, incorporating the supply-side forces highlighted in Katz and Murphy (1992) and Card and Lemieux (2001). These supply terms are measured as the total increase in hours worked per group and instrumented using pre-existing trends in hours during 1970–1980. This strategy isolates the variation in hours due to demographic trends and trends in educational attainment. Controlling for changes in labor supply does not change the qualitative picture, but raises the explanatory power of our task displacement and reinstatement measures. For example, in column 7 Panel C, automation accounts for 66% of variation in between-group wage changes, and new tasks contribute another 37%, while the other variables have a negative contribution. This reflects the fact that demographic trends from 1980 onwards, especially in educational attainment, have gone in favor of groups experiencing more task displacement and less reinstatement during our sample period, and thus, according to our estimated model, without the task displacement and reinstatement developments, these groups would have experienced higher—rather than lower—relative wage growth.

Table 2 turns to analogous specifications for hours worked. Column 1 reports estimates of the bivariate relationship shown in the top panels of Figs. 12 and 13. In Panel A, the coefficient estimate for task displacement is -2.25 (standard error = 0.30), and in Panel B, the coefficient estimate for task reinstatement is 3.62 (standard error = 0.49). Panel C includes both explanatory variables together, with the corresponding coefficients being, respectively, -0.82 (standard error = 0.39) and 2.61 (standard error = 0.71). In this specification, our measures of task changes due to automation and new tasks explain 53% of the variation in changes in hours worked across demographic groups between 1980 and 2016. The remaining columns show that the employment effects are also fairly unchanged when we control for different measures of sectoral reallocation, education and gender dummies, and supply-side factors.

7.5 Robustness

Acemoglu and Restrepo (2022) documented the robustness of the automation results to several other specifications, including those that control for exposure to imports from China and offshoring, those that allow for differential trends for routine jobs and for industries experiencing labor share declines (the two constituent components of our task displacement measure) and those that control for the effects of minimum wages and union coverage. Similar results were also obtained in stacked-differences models and when exploiting variation across US regions.

In the Appendix, we show that the results reported here are robust to the following variations. First, we obtain similar results when we construct the reinstatement due to new tasks using wage-bill variation from 1980 (see Tables A1 and A2). Table A3 decomposes the employment effects into an extensive and intensive margin changes. While the task displacement from automation has

TABLE 2 Reduced-form evidence: changes in hours worked per person regressed on automation and new tasks, 1980–2016.

	Dependent variables: Change in log hours worked per person, 1980–2016						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Panel A. Only displacement from automation						
Automation task displacement	−2.25 (0.30)	−1.58 (0.40)	−1.96 (0.27)	−1.83 (0.40)	−1.93 (0.41)	−2.21 (0.61)	−2.59 (0.78)
R^2 for model	0.44	0.48	0.50	0.68	0.67	0.61	0.56
R^2 for automation	0.44	0.31	0.38	0.36	0.38	0.43	0.51
R^2 remaining covs		0.17	0.11	0.32	0.29	0.18	0.05
Observations	500	500	500	500	500	492	492
	Panel B. Only reinstatement from new tasks						
New tasks reinstatement	3.62 (0.49)	3.40 (0.91)	3.56 (0.46)	1.40 (0.75)	1.46 (0.91)	1.97 (1.19)	3.67 (1.86)
R^2 for model	0.51	0.51	0.51	0.64	0.62	0.22	−0.09
R^2 for new tasks	0.51	0.48	0.50	0.20	0.20	0.28	0.51
R^2 remaining covs		0.03	0.01	0.44	0.41	−0.06	−0.61
Observations	500	500	500	500	500	492	492
	Panel C. Both explanatory variables						
Automation task displacement	−0.82 (0.39)	−0.81 (0.40)	−0.95 (0.40)	−1.75 (0.40)	−1.86 (0.40)	−2.13 (0.59)	−2.55 (0.77)
New tasks reinstatement	2.61 (0.71)	2.48 (0.95)	2.16 (0.61)	0.58 (0.63)	0.93 (0.79)	0.61 (0.68)	1.55 (0.85)
R^2 for model	0.53	0.53	0.53	0.68	0.67	0.61	0.55
R^2 for automation	0.16	0.16	0.19	0.34	0.37	0.42	0.50
R^2 for new tasks	0.37	0.35	0.30	0.08	0.13	0.08	0.22
R^2 remaining covs		0.02	0.04	0.26	0.17	0.11	−0.17
R^2 remaining covs	500	500	500	500	500	492	492
	Panel D. Net task change due to new tasks minus automation						
Net task change (new tasks−automation)	1.52 (0.19)	1.32 (0.32)	1.37 (0.17)	1.41 (0.30)	1.63 (0.33)	1.76 (0.49)	2.39 (0.71)
R^2 for model	0.51	0.52	0.53	0.68	0.67	0.58	0.53
R^2 for net task changes	0.51	0.45	0.46	0.48	0.55	0.59	0.81
R^2 remaining covs		0.07	0.06	0.20	0.12	−0.01	−0.28
Observations	500	500	500	500	500	492	492
<i>Other covariates:</i>							
Sectoral value added		✓		✓		✓	
Sectoral TFP			✓		✓		✓
Sectoral markups				✓	✓		✓
Gender and education dummies				✓	✓	✓	✓
Labor supply shifts					✓		✓

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hours worked per person across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The dependent variable is the change in log hours per person for each group between 1980 and 2016. Panel A reports results using only our task displacement measure. Panel B only uses our task reinstatement measure. Panel C includes both task displacement and task reinstatement on the right-hand side. Panel D combines task displacement and reinstatement into a single net task change measure. The bottom rows list additional covariates included in each specification. As in Acemoglu and Restrepo (2022), we instrument changes in labor supply in columns 6 and 7 using trends in total hours worked by group from 1970 to 1980. All regressions are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

a robust negative association with both margins, new tasks are more strongly associated with increases in employment at the extensive margin. Finally, we show in Table A4 that the coefficients on task displacement and reinstatement variables are comparable when we estimate the models separately for workers with and without a college degree. This exercise shows that the benefits from new tasks and the costs of automation are visible even when focusing on these specific segments of the labor force.

7.6 Taking stock

Our reduced-form findings support the main implications of the task framework: task displacement due to automation has a sizable negative effect on the relative wages of exposed groups, and reinstatement driven by new tasks has a sizable positive effect on relative wages. These two variables explain at least 60% of the total variation in between-group wage changes between 1980 and 2016. Consistent with the expectation that these technology measures shift the relative demand for labor from different skill groups, we find that they have commensurate effects on employment as well. The two measures together account for approximately 53% of the variation in the changes in hours worked for the same time period. In line with our theory, the estimates also suggest that technologies that cause extensive-margin changes (thus reallocating tasks from one factor to another) explain the bulk of variation in the changes in the wage and employment structure, and have much greater explanatory power than proxies for factor-augmenting and sectoral technology variables.

Despite the clear empirical associations uncovered here, it is important to exercise caution in interpreting these reduced-form results. First, our proxies for factor-augmenting and sectoral changes are imperfect. The education dummies may capture other trends as well as factor-augmenting technologies, while the reduced-form estimates of the contribution of sectoral variables may be attenuated. Second, we are ignoring ripple effects, which link the wages of a skill group to the task displacement experienced by other groups of workers—especially when there are high levels of substitutability between the groups in question. Third, productivity effects are subsumed into the constant. All of these considerations motivate our approach in the next section, which further leverages the structure of the model to estimate the propagation matrix and productivity implications of different types of technologies, and performs counterfactual exercises to measure their contribution to the changes in wage inequality since 1980.

8 Estimation of general equilibrium effects and counterfactuals

This section uses the task model to quantify the equilibrium effects of different technologies on the US wage structure. We use the equations characterizing

the impacts of technology, inclusive of the ripple effects. We implement these equations using the measures of the direct effects of different technologies introduced above, and combine them with external information on a number of key elasticities of substitution and our estimates of the propagation matrix.

This exercise adds to the reduced-form findings in three ways. First, it accounts for the effects of technology on wage levels working via the productivity effect. As explained previously, the reduced-form evidence is informative about relative change in wages and employment for exposed groups—but not about the effects of different technologies. Second, it enables us to estimate ripple effects. Finally, this exercise incorporates the effects of technology working through changes in sectoral composition. Reduced-form models controlled for sectoral shifts but did not estimate the effects of different types of technologies working through the sectoral changes that they induced. Our results from this structural exercise suggest that automation and new tasks are important drivers of the changes in the US wage structure.

8.1 General equilibrium effects of technology and markups

Our objective is to estimate separate effects of automation, new tasks, Hicks-neutral sectoral productivity (TFP) shifters and markups on hourly wages. We return to the contribution of factor-augmenting technologies later. The analysis can be expanded to include other factors, but we do not do so to keep the chapter focused on the consequences of technology trends.

From Propositions 9 and 10, the change in group wages can be written as

$$\begin{aligned} d \ln w = \Theta \cdot \text{stack} \left(& d \ln y - d \ln M - d \ln \Gamma_g^{\text{auto}} + d \ln \Gamma_g^{\text{new}} \right. \\ & - (1 - \lambda) \cdot \sum_i \omega_{gi} \cdot d \ln A_i - \lambda \cdot \sum_i \omega_{gi} \cdot d \ln \mu_i \\ & \left. + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right) + \nu. \end{aligned}$$

In this equation, ν is an error term subsuming all other forces shaping the wage structure. The endogenous price changes $\{d \ln p_i\}_i$ associated with these shocks satisfy

$$\begin{aligned} d \ln p_i = \sum_g s_g^{y_i} \cdot d \ln w_g - \sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{auto}} \cdot \pi_{gi}^{\text{auto}} - \sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{new}} \cdot \pi_{gi}^{\text{new}} \\ - d \ln A_i + d \ln \mu_i. \end{aligned} \tag{41}$$

To determine the effects of these technologies on output, we simplify the analysis by assuming that, initially, $\mu_i = 1$ for all i . This assumption implies that the sectoral value-added shares are equal to sectoral cost shares and that, as in

Section 5, the change in aggregate output, $d \ln y$, is determined by the following equation, which relates average wage changes to changes in TFP and markups:

$$\sum_g s_g^y \cdot d \ln w_g = \sum_i s_i \cdot \left[\sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{auto}} \cdot \pi_{gi}^{\text{auto}} \right. \\ \left. + \sum_g s_g^{y_i} \cdot d \ln \Gamma_{gi}^{\text{new}} \cdot \pi_{gi}^{\text{new}} + d \ln A_i - d \ln \mu_i \right]. \quad (42)$$

Because we are looking at first-order approximations, these three equations provide an additive decomposition of the contribution of technologies and markups.

To implement this decomposition, we need estimates of (i) initial factor shares; (ii) elasticities $\{\lambda, \eta\}$; (iii) direct task displacement and reinstatement, $\{d \ln \Gamma_g^{\text{auto}}, d \ln \Gamma_g^{\text{new}}\}_g$; (iv) sectoral TFP growth, $\{d \ln A_i\}_i$, and sectoral markup changes $\{d \ln \mu_i\}_i$; and (v) the propagation matrix, Θ .

- For (i), we take factor shares directly from the Census data matched to the BEA-BLS industry accounts.
- For (ii), we set $\lambda = 0.5$ and $\eta = 0.3$. The task-elasticity of substitution λ comes from Humlum (2020), who estimates it on Danish manufacturing data. The estimate for the sectoral elasticity of substitution is from Buera et al. (2021) and is a standard value used in the structural transformation literature.
- For (iii), we continue to use the measure of new task reinstatement in (40), but a slightly different measure for task displacement due to automation, given by

$$d \ln \Gamma_{gi}^{\text{auto}} = \text{RCA routine}_{gi} \cdot \frac{-\Delta \ln s_L^{y_i, \text{auto}}}{1 + (\lambda - 1) \cdot s_\ell^{y_i} \cdot \pi_i^{\text{auto}}} \quad (43)$$

for group g in industry i . This expression differs from the measure used in the reduced-form analysis, Eq. (38), because of the term $(\lambda - 1) \cdot s_\ell^{y_i} \cdot \pi_i^{\text{auto}}$ in the denominator, which adjusts for the effect of automation on the labor share working via substitution towards the cheaper newly-automated tasks. The earlier expression obtains when $\lambda = 1$. We used this restriction in our reduced-form analysis to simplify the exposition. Here, we construct the adjustment term using $\lambda = 0.5$ and $\pi_i^{\text{auto}} = 30\%$. Total task displacement due to automation $d \ln \Gamma_g^{\text{auto}}$ aggregates the new measures for $d \ln \Gamma_{gi}^{\text{auto}}$ across industries, as in Eq. (39).³¹

To obtain cost savings from these technologies, we follow Acemoglu and Restrepo (2022) and set $\pi_{gi} = 30\%$. This choice is motivated by available estimates of cost savings due to the adoption of industrial robots in US manufacturing. This choice assumes the same savings for automation in other

³¹ The reduced-form results are very similar with the adjusted measure shown here and other variants, and are presented in the Appendix of Acemoglu and Restrepo (2022).

sectors, which is an assumption that can be relaxed in the future using additional data. For new tasks, we set $\pi_{gi}^{new} = 30\%$ for symmetry, since we do not have direct estimates of the surplus generated by new tasks. This number implies that a 10% increase in new tasks for all worker groups would raise TFP by 3%, which is a reasonable number.³²

- For (iv), we estimate the sectoral Hicks-neutral productivity shifters $\{d \ln A_i\}_i$'s by subtracting the implied TFP gains due to automation and new tasks from observed industry TFP changes. The left panel in Fig. 14 depicts observed industry TFP changes together with the implied estimates for the $d \ln A_i$'s.³³ Computers and electronics and transportation pipelines experienced the largest sectoral productivity increases, while legal services and transportation services experienced the least. Overall, the two series are highly correlated, but there are some notable exceptions, such as motor vehicles, where observed TFP exceeds our estimate for $d \ln A_i$ by a sizable amount, since this industry has made large automation investments during this period.

For markups, we use the estimates from Hubmer and Restrepo (2021). These are estimated using the production function approach and Compustat data as in De Loecker et al. (2020), but allow firm-level output elasticities to vary by size, and also aggregate these markups using their sales-weighted harmonic mean to obtain aggregate industry markups. These estimates are shown in the right Panel of Fig. 14.

8.2 Estimating the propagation matrix

The wage equation in the multi-sector model, (28), can be rewritten as

$$\Delta \ln w_g = \frac{1}{\lambda} \cdot (d \ln \Gamma_g^{new} - d \ln \Gamma_g^{auto}) + \beta \cdot X_g + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g}{\partial \ln w} \cdot \text{stack}(\Delta \ln w_{g'}) + u_g, \quad (44)$$

where X_g is a vector that contains sectoral shifts and education and gender dummies, proxying for other technological trends. Rather than solving out for the vector of wage effects using the propagation matrix as in (28), here we include the vector of wage changes for other demographic groups on the right-hand side, which highlights that these will impact the wage of group g via the g th row of the task-shares Jacobian matrix, $\frac{\partial \ln \Gamma}{\partial \ln w}$. The error term u_g contains all unobserved labor demand and supply shocks impacting demographic group g .

³² Our prior is that this number should be bigger since new tasks enable various efficiency-enhancing improvements and the reorganization of production process as explained in footnote 14. Nevertheless, we choose 30% to err on the conservative side.

³³ For simplicity, our theory used value-added production functions at the industry level (with material inputs solved out). To match this choice, we use measures of value-added TFP instead of gross-output TFP. While it would be preferable to use measures of TFP for gross output (so that they can be readily interpreted as technology), this would require modeling input-output linkages across industries, which we do not pursue for this chapter.

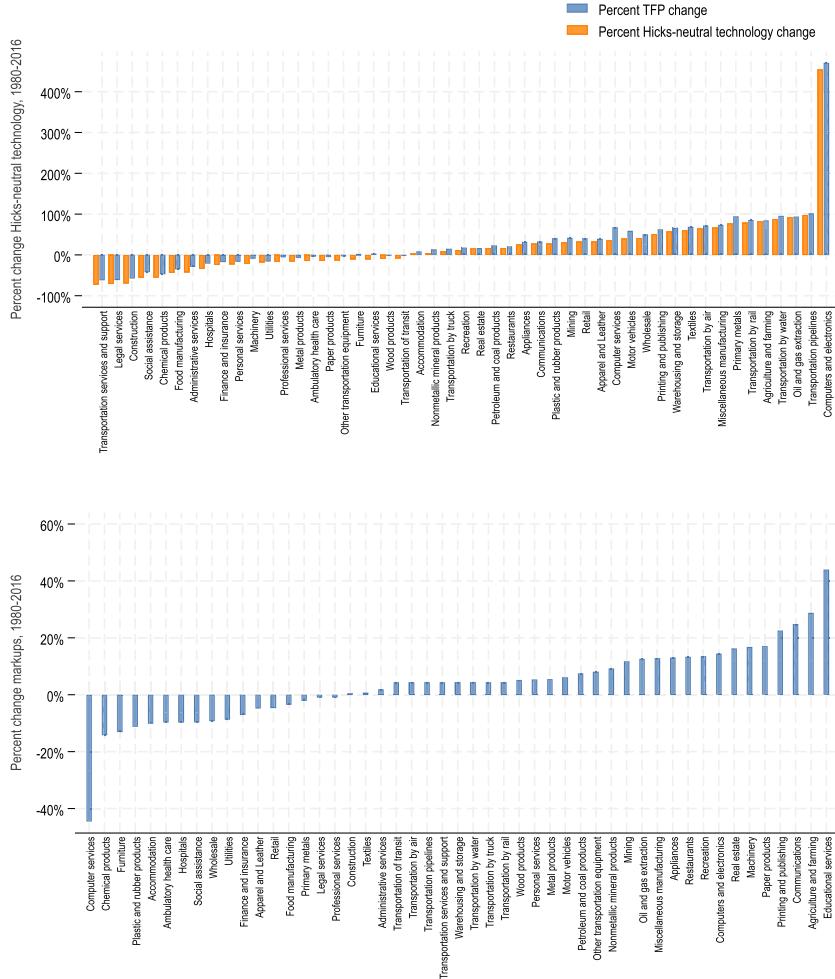


FIGURE 14 The top panel depicts the percent change in TFP (in blue) and Hicks-neutral technology (in orange) for US industries. The Hicks-neutral component is obtained by subtracting the contribution of new tasks and automation to sectoral TFP. The bottom panel provides estimates for the change in markups across US industries, from Hubmer and Restrepo (2021).

Our strategy is to estimate the Jacobian using GMM (Generalized Method of Moments). In this estimation, we impose external values for λ and use the orthogonality conditions

$$d \ln \Gamma_g^{\text{auto}}, d \ln \Gamma_g^{\text{new}}, X_g \perp u_{g'} \text{ for } g, g' \in \mathbb{G},$$

which impose that task displacement and reinstatement terms as well as the education and gender dummies and sectoral shifters in X_g are orthogonal to the error term. This orthogonality assumption was implicit in the reduced-form

models estimated in the previous section. Once the Jacobian matrix is estimated, the propagation matrix can be obtained as $\Theta = \frac{1}{\lambda} \cdot \left(\mathbb{1} - \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma}{\partial \ln w} \right)^{-1}$.

The Jacobian is a $G \times G$ matrix, and hence it would be impossible to estimate all of its entries in an unrestricted fashion. Instead, we follow Acemoglu and Restrepo (2024) and parameterize the entries of the Jacobian in terms of similarities between groups.³⁴ This approach operationalizes the intuitive idea that the Jacobian matrix is informative about the extent of substitutability between groups and such substitutability should depend on how similar the groups are. We assume that the off-diagonal terms of the Jacobian (for $g' \neq g$) can be parameterized as

$$\frac{\partial \ln \Gamma_g}{\partial \ln w_{g'}} = s_{g'}^y \cdot \varphi + \sum_n \omega_{gn} \cdot s_{g'}^n \cdot [\gamma + \gamma_{\text{job}} \cdot \text{job similarity}_{gg'} + \gamma_{\text{edu-age}} \cdot \text{edu-age similarity}_{gg'}],$$

while the diagonal terms take the form

$$\frac{\partial \ln \Gamma_g}{\partial \ln w_g} = (s_g^y - 1) \cdot \varphi - \sum_n \sum_{g' \neq g} \omega_{gn} \cdot s_g^n \cdot [\gamma + \gamma_{\text{job}} \cdot \text{job similarity}_{gg'} + \gamma_{\text{edu-age}} \cdot \text{edu-age similarity}_{gg'}].$$

This parameterization implies that competition for marginal tasks between skill groups takes place within job categories, denoted by n . In the data, we assume that there are 96 job categories, given by combinations of 16 aggregated industries and six aggregated occupations. The summation terms indicate that the effects of competition from group g' on group g in category n depend on the importance of this category for group g , summarized by the share of category n in the total wage payments for group g (ω_{gn}), and the share of wage payments in job category n accruing to group g' ($s_{g'}^n$). Both of these objects are computed from the 1980 Census. Intuitively, groups with greater wage shares should generate more competitive pressure on other groups in the same job category, as implied, for example, by the Frechet parameterization of comparative advantage in Section 4. In addition, the three terms in square brackets represent three dimensions of competition between groups. The first, with coefficient $\gamma \geq 0$, corresponds to the component of competition that is common to all workers in a job category. The second, with coefficient $\gamma_{\text{job}} \geq 0$, is from the similarity of the jobs performed by the two demographic groups. In particular, we use the *cosine* similarity of job categories performed by groups g' and g in the 1980 Census. This functional form is also motivated by the Frechet example, where a higher

³⁴ In Acemoglu and Restrepo (2022), we directly parameterized and estimated the propagation matrix. We prefer the current approach because it is easier to develop an intuition about the entries of the Jacobian, which correspond to first-round ripple effects (rather than the Leontief inverse of this matrix, which depends on higher-round ripples).

correlation in task-level productivities results in higher substitutability. The third term, with coefficient $\gamma_{\text{edu-age}} \geq 0$, parameterizes the extent to which competition for tasks is stronger for workers of similar education and experience, as in Card and Lemieux (2001). We compute this similarity measure as follows: we run a Mincer wage equation for log hourly wages in 1980, as a function of age and education dummies, and then construct the education-age similarity between two groups as the inverse distance between the predicted wage level of groups g and g' in 1980. This procedure captures how similar the two groups are in terms of their education and age, with each of these dimensions weighted by their Mincer coefficients.

Finally, the parameter $\varphi \geq 0$ regulates the extent of competition between capital and workers for marginal tasks, which is assumed to be uniform across groups. Our parametrization implies that the row sums of the Jacobian are equal to $-s_K \cdot \varphi$. Using the definitions in Section 4, we see that the macroeconomic elasticity of substitution between capital and labor is $\sigma_k = \lambda + \varphi$.³⁵ We set $\varphi = 0.1$, so that σ_k matches estimates of the elasticity of substitution between capital and labor in Oberfield and Raval (2021) of around 0.6. This parameterization therefore fixes the row sums of the Jacobian, $\frac{\partial \ln \Gamma}{\partial \ln w}$, and allows the data to determine the γ coefficients, which determine the strength of competition for marginal tasks between different groups.

Table 3 reports our estimates for the γ 's obtained from Eq. (44). For these estimates, we additionally impose the restriction that $\gamma, \gamma_{\text{job}}, \gamma_{\text{edu-age}} \geq 0$. When we include all three terms simultaneously, the first two are estimated to have zero coefficients (given our nonnegativity constraint) and the spillover patterns are explained by the education-age similarity measure, as in the specifications in columns 3 and 6. In what follows, we take column 3—which has $\gamma = 0$, $\gamma_{\text{job}} = 0$, and $\gamma_{\text{edu-age}} = 0.8$ —as our preferred specification.

The estimated propagation matrix has an average diagonal of 0.84, and the row sum of the off-diagonal terms is about 1. This implies that workers from group g bear about 45% of the incidence of a direct shock reducing their labor demand, with the rest being shifted to other groups via competition for marginal tasks.

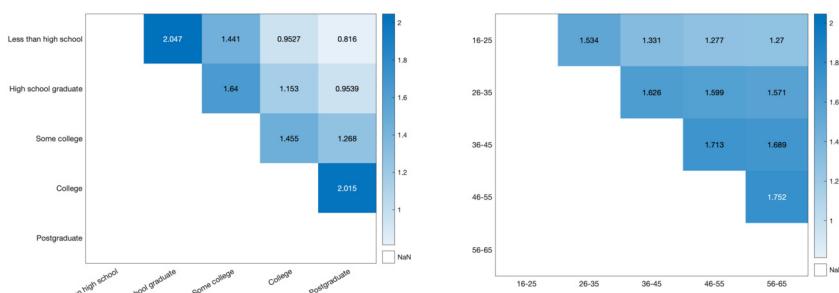
Another way to illustrate the structure of the estimated propagation matrix is by looking at the implied elasticity of substitution between skill groups. Fig. 15 provides this information by computing the unweighted average of pairwise elasticities of substitution across indication groups (on the left) and age groups (on the right). The average elasticity of substitution between groups with a college and postgraduate degree is estimated to be 2, while the average elasticity of substitution between groups with a college degree and those without a high school degree is 0.95.

³⁵ Recall that due to symmetry, $\sigma_{kg} = \sigma_{gk}$. Moreover, we can write $\sigma_{gk} = \lambda + \frac{1}{s_K^3} \left(-\sum_{g'} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}} \right)$, since a change in the cost of capital is equivalent to an increase in all wages. This implies $\sigma_{kg} = \sigma_{gk} = \lambda + \varphi$.

TABLE 3 Estimates of the task-shares Jacobian.

	Dependent variables: Change in log hourly wages, 1980–2016					
	(1)	(2)	(3)	(4)	(5)	(6)
Baseline competition γ	0.39 (0.13)			0.33 (0.16)		
Job-similarity competition γ_{job}		0.74 (0.21)			0.68 (0.25)	
Education-age competition $\gamma_{\text{edu-age}}$			0.80 (0.22)			0.84 (0.31)
Observations	500	500	500	500	500	500
Covariates:						
Gender and education dummies	✓	✓	✓	✓	✓	✓
Sectoral value added control	✓	✓	✓		✓	✓
Sectoral TFP and markups				✓	✓	✓

Notes: This table presents estimates of the (task share) Jacobian, using the parameterization in Section 8. The estimation equation can be written as $\sigma \Delta \ln w_g + d \ln \Gamma_g^{\text{auto}} - d \ln \Gamma_g^{\text{new}} = \tilde{\beta} X_g + \gamma \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) + \gamma_{\text{job}} \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{job similarity}_{gg'} \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) + \gamma_{\text{edu-age}} \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{edu-age similarity}_{gg'} \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) + \tilde{v}$, where $\tilde{\beta}$ and \tilde{v} are linear transformations of β and v respectively. The ripple terms are instrumented using $\sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g)$, $\sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{job similarity}_{gg'} \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g)$ and $\sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{edu-age similarity}_{gg'} \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g)$, respectively, where $\Delta \ln \hat{w}_g$ is the predicted wage change based on task displacement, task reinstatement and the covariates. Columns 1 and 4 present estimates for γ excluding the other two spillover terms. Columns 2 and 5 present estimates for γ_{job} excluding the other two spillover terms. Columns 3 and 6 present estimates for $\gamma_{\text{edu-age}}$ excluding the other two spillover terms. When all three measures of competition are included and the restriction that they must have non-negative coefficient is imposed, the first two are estimated to have zero effects and the results are identical to those in columns 3 and 6. All estimates are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

**FIGURE 15** The figure reports average elasticities of substitution between educational and age groups. These averages are obtained from our estimates of the propagation matrix.

8.3 Decompositions

We first illustrate the effects of each type of technological change, highlighting the different pathways via which they affect labor demand.

Fig. 16 depicts the effects of automation. The panels plot estimates of the different mechanisms, which we accumulate from left to right, with the rightmost panel corresponding to the total effect of the technology in question. The vertical axes show the model estimates (in units of change in hourly wages from 1980 to 2016), while the horizontal axis ranks groups according to hourly wages in 1980. Panel A starts with the productivity gains from automation, $(1/\lambda) \cdot d \ln y$. We see here that automation increased output by 20% over this period, which raised the wages by 40%.

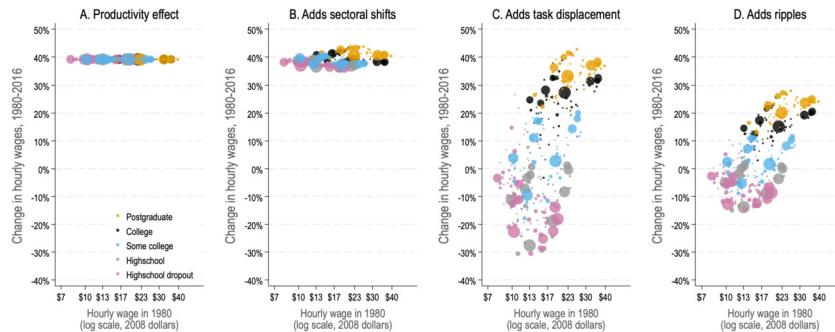


FIGURE 16 This figure decomposes the effects of automation on hourly wages between 1980 and 2016 into four components. Panels sequentially add productivity effects, industry shifts, task displacement from automation, and ripple effects. The horizontal axis ranks groups according to hourly wages in 1980. Marker sizes are proportional to hours worked in 1980, and marker colors distinguish groups by education.

Panel B adds the effects of automation working through changes in the sectoral composition of the economy by plotting $(1/\lambda) \cdot (d \ln y + (\lambda - \eta) \sum_i \omega_{gi} \cdot d \ln p_i)$. Note that here we only account for the change in sectoral prices due to automation, computed according to Eq. (41). The change in sectoral prices due to automation does not generate much variation in terms of relative wage changes. This is because the skill composition of sectors expanding due to automation is similar to the rest.

Panel C adds the direct task displacement due to automation and plots $(1/\lambda) \cdot (d \ln y - d \ln \Gamma_g^{\text{auto}} + (\lambda - \eta) \sum_i \omega_{gi} \cdot d \ln p_i)$. The uneven impacts across groups are now clearly visible. For example, task displacement reduces the wages for some groups by as much as 30%, while the real wages of highly-educated groups shielded from automation increase by more than 40%. This panel confirms that automation works primarily by displacing workers from their tasks, shifting labor demand within sectors—rather than by shifting the sectoral composition of the economy, as in Panel B.

Panel D adds the ripple effects generated by automation. We see here that ripples play an equalizing role, consistent with our discussion in Section 5. This

is because groups that experience a large reduction in their task share due to automation are able to compete for marginal tasks previously performed by other groups. This reallocation spreads the negative incidence of automation to other groups and mitigates the adverse effects on exposed groups. Our estimates imply that high school graduates experienced on average a 4.3% wage decline due to automation, and groups with less than high school experienced even steeper declines of 8.1%. College graduates and postgraduates, on the other hand, enjoyed 17.6% and 22.9% wage increases from automation. Underscoring the equalizing role of the ripple effects, the declines in the real wages of high school graduates and less than high school groups would have been, respectively 10.1% and 16.2%, if these groups had not been able to compete for marginal tasks and shift some of the burden of task displacement to other groups.

Fig. 17 depicts the estimated effects of new tasks on wages from 1980 to 2016. The panels have the same organization as before. Our estimates imply that new tasks reduce output by a small amount. This does not mean that the economy is made less productive by new tasks. In fact, new tasks raise TFP by 5%, and average wages and aggregate consumption by 7%. The reason why output declines is because new tasks make the production process less capital intensive and as a result the share of capital and investment decrease (recall the relationship between TFP change and output change in footnote 10).

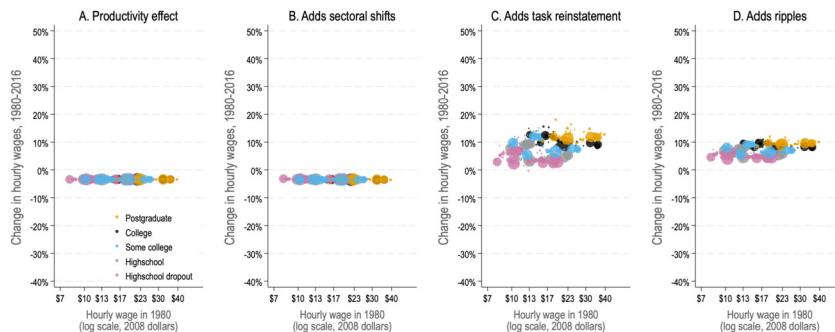


FIGURE 17 This figure decomposes the effects of new tasks on hourly wages between 1980 and 2016 into four components. Panels sequentially add productivity effects, industry shifts, task reinstatement from new tasks, and ripple effects. The horizontal axis ranks groups according to hourly wages in 1980. Marker sizes are proportional to hours worked in 1980, and marker colors distinguish groups by education.

New tasks benefit all groups but generate more pronounced gains for highly-educated and highly-paid workers. New tasks thus contributed to rising inequality, even if by a much smaller amount than automation. This result aligns with our reduced-form findings, where automation explains a larger share of the observed variance in wage trends than do new tasks. The overall wage increase due to new tasks ranges from 5.30% for groups with less than high school to

10.1% for college workers in Panel C. This heterogeneity is, as usual, further compressed by the ripple effects in Panel D.³⁶

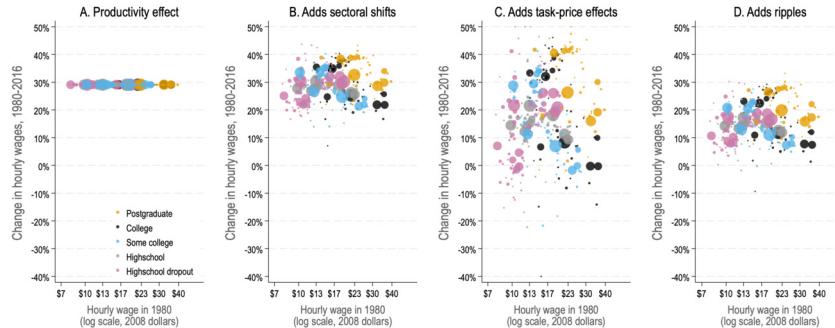


FIGURE 18 This figure decomposes the effects of sectoral TFP changes on hourly wages between 1980 and 2016 into four components. Panels sequentially add productivity effects, industry shifts, effects via task prices, and ripple effects. The horizontal axis ranks groups according to hourly wages in 1980. Marker sizes are proportional to hours worked in 1980, and marker colors distinguish groups by education.

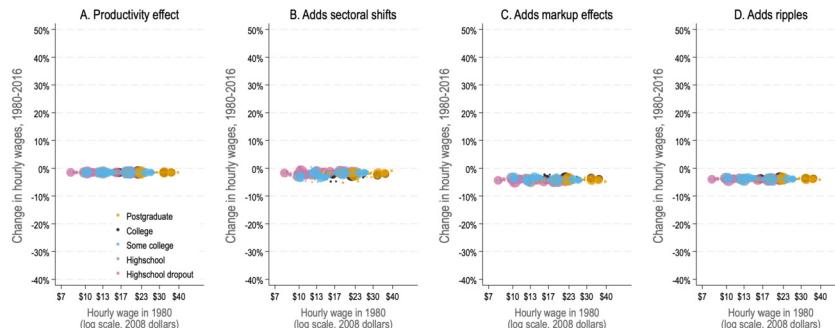


FIGURE 19 This figure decomposes the effects of sectoral markups on hourly wages between 1980 and 2016 into four components. Panels sequentially add productivity effects, industry shifts, direct effects of markups, and ripple effects. The horizontal axis ranks groups according to hourly wages in 1980. Marker sizes are proportional to hours worked in 1980, and marker colors distinguish groups by education.

Figs. 18 and 19 plot the results for sectoral TFP changes and markups, which are estimated to have modest distributional implications. Changes in sectoral TFP increase wages for all groups by about 15%. Due to the fact that $\eta < 1$, they also reallocate labor towards high-skill services, which benefits workers with a post-graduate degree.³⁷

³⁶ New tasks increase the total mass of tasks M by $d \ln M = (1 - (\lambda - 1) \cdot \pi_g^{new}) \cdot \sum_{g \in \mathcal{G}} \pi_g^L \cdot d \ln \Gamma_g^{new}$. This effect is common to all workers and is included in Panel C.

³⁷ This is in line with previous work by Buera et al. (2021), who also document that the process of structural transformation in the US raised the relative demand for college-educated workers.

Markups reduce output and real wages, but affect groups uniformly. This is because the sectors experiencing the most pronounced increase in markups are similar to the rest in terms of the composition of their workforces.

Fig. 20 aggregates the effects of automation, new tasks, sectoral TFP changes, and markups for 1980-2016 and compares their estimated wage impacts to observed wage changes in this period. These trends combined account for 72% of between-group wage changes from 1980 to 2016.

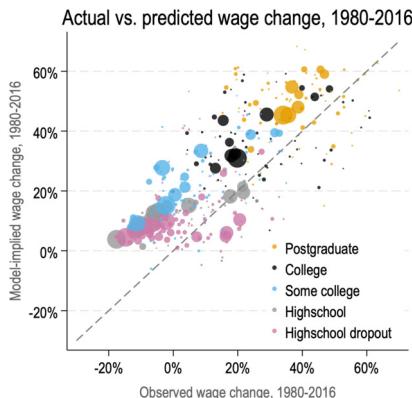


FIGURE 20 The figure plots observed wage changes in (real) hourly wages, 1980-2016, vs. predicted changes based on the combined effects of automation, new tasks, sectoral TFP changes, and sectoral markup changes estimated using our model.

Table 4 summarizes the individual contribution of the different technologies studied here and sectoral markups to the observed wage changes. Automation technologies introduced since 1980 account for 55% of the observed wage trends across worker groups. New tasks contributed 8.7%, as they have favored highly-educated workers the most. Changes in sectoral TFP contributed 7.5%, while changes in sectoral markups had minor effects.

The second column reports predicted average wage growth coming from each source. Despite generating large distributional effects, automation brought a modest increase in average wages of about 4.4%. The opposite holds for Hicks' neutral sectoral TFP improvements, which increased average wages by 15.4%, with modest distributional effects in comparison. Overall, predicted wage growth from the model exceeds the composition-adjusted real wage growth in the US economy over the same time period, which is about 5%. This could be because other factors (for example, related to non-competitive elements in the labor market discussed below) may have put additional downward pressure on wages.

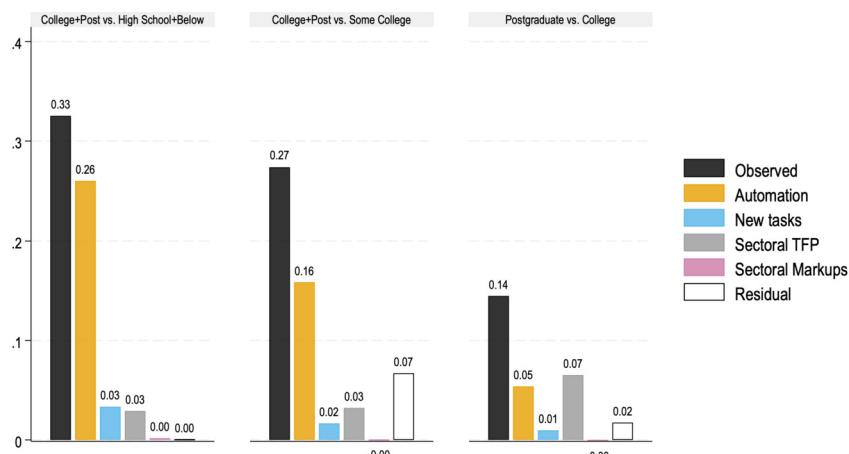
Fig. 21 provides additional details on the impacts of different types of technologies on the wage structure. It depicts the contribution of the same four factors to the wage premium earned by college-educated workers relative to those with high school or less; the premium of college-educated workers relative

TABLE 4 Share of variance in wage trends across groups explained by different technologies and markups.

	Share wage changes explained, 1980–2016 (1)	Contribution to average wage growth, 1980–2016 (2)
Automation	55.34%	4.38%
New task creation	8.70%	7.06%
Sectoral TFP changes	7.47%	15.39%
Markups	0.69%	−3.87%
Total	72.20%	22.96%

Notes: Column 1 reports the contribution of the indicated technology term to observed wage changes across 500 demographic groups between 1980 and 2016. This is weighted by total hours worked by each group in 1980. Column 2 reports predicted average (real) wage growth between 1980 and 2016 from the indicated types of technological change.

to those with some college; and the premium earned by postgraduate workers relative to those with a college degree. Automation is the most important driver of the increase in the college premium and also plays an important role in explaining the rising postgraduate premium. New tasks and sectoral TFP trends also contributed to the rising college premium, though with a smaller role than automation. Sectoral TFP trends had a more prominent role in explaining the rise in the postgraduate premium since 1980, partly because a few sectors that disproportionately employ postgraduates, such as legal services and health care, experienced lackluster productivity growth, which led to their expansion as a share of value added.

**FIGURE 21** The figure reports the estimated contribution of technology and markups to the changes in various educational premia, 1980–2016. The bars represent the effects of different technologies or sectoral markup changes.

8.4 Limited distributional impacts of labor-augmenting technologies

Our decomposition exercise ignored the role of labor-augmenting technological changes, because we have no direct measures of such technologies. In this subsection, we perform a bounding exercise to show these technologies are unlikely to be important drivers of the changes in the US wage structure between 1980 and 2016.

We consider three types of technological changes: automation, uniformly labor-augmenting technologies, and labor-augmenting technologies at the intensive margin. For each technology, we consider a shock that generates a 1% increase in TFP and then trace its contribution to inequality. Because each of the shocks we are considering is chosen to raise TFP by 1%, we know from theory that their impact on average wages is to increase them by 1.5% (this follows from $\sum_g s_g^y \cdot d \ln w_g = d \ln tfp$).

In the top panel of Table 5, we investigate how large the distributional effects of automation are relative to their TFP impact. We consider advances in automation equally affecting all skill groups with the same education level. For example, the first row considers the hypothetical effects of advances in automation affecting only high-school dropouts, and reports the effects of these advances on the wages of workers of different educational levels (in each case averaged across demographic groups with the same level of education). In this exercise we keep π_g^{auto} fixed and set the fraction of automated tasks to ensure a 1% increase in aggregate TFP.

Panel A shows that automation has significant distributional effects. For instance, a (uniform) automation shock impacting all groups with less than high school reduces these groups' own wage by, on average, -21.88%. The impact on other demographic groups, operating via the productivity and ripple effects, is positive. For instance, the effect on college-graduate groups is an 8.07% increase. This implies that automation affecting workers with less than high school is increasing inequality between this group and college graduates by about 30%.

Panel B shows positive but comparatively much smaller effects on own group wages from uniformly labor-augmenting technologies, which reflects the fact that the macroeconomic elasticities between groups (taking into account the ripple effects) are close to 1. For example, a technological improvement raising the productivity of workers with less than high school degree uniformly increases their wages by 2%, and has a very similar impact on groups with college or more. The quantitative pattern in the other rows is similar: uniformly labor-augmenting technologies have a limited effect on inequality and generate similar wage gains across all educational groups.

Panel C of Table 5 repeats this exercise for labor-augmenting changes at the intensive margin. As highlighted in Proposition 6, these technologies have a more negative impact on the group experiencing the increase in productivity because they do not generate the same beneficial impact via competition for marginal tasks. This is why the diagonal in Panel C with the own-group effects is negative. Despite reducing the wages of exposed groups, the effects of this form

TABLE 5 Effects on average wages due to a 1% increase in TFP by demographic group.

Shock to	Effects on average real hourly wages (%):				
	High School Dropout	High School Graduate	Some College	College	Postgraduate
Panel A. Automation					
High School Dropout	−21.88	4.25	5.98	8.07	8.86
High School Graduate	4.15	−8.57	5.38	7.06	8.05
Some College	5.8	5.26	−13.41	5.9	6.68
College	7.96	6.92	5.71	−27.65	3.13
Postgraduate	9.12	8.15	6.47	2.63	−27.86
Panel B. Uniform factor-augmenting					
High School Dropout	2.02	0.84	1.34	1.94	2.16
High School Graduate	0.81	1.85	1.16	1.65	1.93
Some College	1.29	1.13	2.29	1.32	1.54
College	1.91	1.61	1.26	2.30	0.51
Postgraduate	2.24	1.96	1.48	0.36	0.84
Panel C. Intensive-margin factor-augmenting					
High School Dropout	−2.11	1.88	2.20	2.56	2.75
High School Graduate	1.85	−0.03	2.09	2.38	2.58
Some College	2.16	2.06	−0.74	2.17	2.32
College	2.54	2.35	2.14	−2.91	1.67
Postgraduate	2.78	2.59	2.27	1.60	−3.16

Notes: This table shows the effects on average wages in demographic groups due to a rise in factor-augmenting technologies that result in a 1% increase in TFP. The detailed breakdown by panel facilitates understanding of the differential impact across various scenarios of technological advancement and educational strata.

of technology on inequality are modest, especially when compared to the effects of automation in Panel A. For example, an intensive-margin labor-augmenting technology raising the productivity of skill groups with less than high school reduces their wages by about 2.11% and increases the wages of other groups by 1.88%–2.75%, thus amounting to a 4.5% widening of between-group wages. This quantitative impact is an order of magnitude smaller than the distributional implications of automation technologies.

The limited distributional impacts of labor-augmenting technologies are also implied by the small explanatory power of the education and gender dummies estimated in the reduced-form models, recalling that these flexibly subsume education-augmenting and gender-augmenting technological developments. Overall, factor-augmenting technologies appear to have fairly limited distributional effects in this framework.

We have so far emphasized the success of the task framework in accounting for various recent labor market trends. We conclude this section by highlighting two puzzles that this framework generates, which require further work.

8.5 The missing technology puzzle

Our decomposition exercise focused on accounting for wage changes across skill groups. A related but distinct exercise is to explore the contribution of dif-

ferent technological trends to total demand shifts. Since 1980, the US workforce has become significantly more educated, which translates into large changes in the size of more educated skill groups. As emphasized in Katz and Murphy (1992), all else equal, this demographic shift should have raised the relative wages of less educated workers. From the viewpoint of the standard relative supply-demand framework, this implies that the relative demand changes have been even larger and have favored the more educated groups.

Following Katz and Murphy (1992), we can use the framework here to quantify the extent of these demand shifts. In particular, given the propagation matrix Θ , which summarizes all the relevant elasticities, the demand shifts across demographic groups since 1980 can be computed as

$$\text{demand shift}_g = \Delta \ln w_g + \Theta_g \cdot \text{stack}(\Delta \ln \text{population}_g + \Delta \ln \ell_g), \quad (45)$$

where $\Delta \ln \text{population}_g$ are changes in log group size and $\Delta \ln \ell_g$ denotes changes in log hours per capita. This expression leverages the fact that the propagation matrix also controls how changes in the supply of skills affect wages, as explained in Proposition 7.

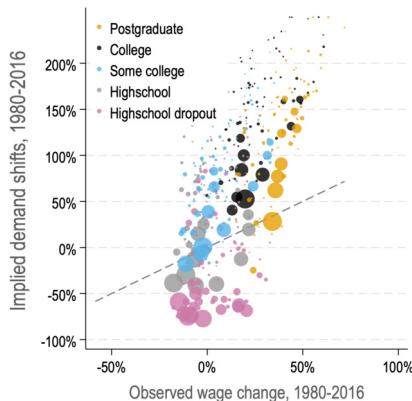


FIGURE 22 The figure plots the total demand shifts computed from Eq. (45) between 1980 and 2016 for 500 demographic groups. These are compared to observed wage changes during this period in the horizontal axis. Marker sizes are proportional to hours worked in 1980, and marker colors indicate education levels.

Fig. 22 compares the measured demand shifts with observed wage changes and underscores the point we made above: demand shifts are more pronounced than wage movements because supply shifts have favored low-education and low-pay groups. But then, what explains these demand shifts? According to our estimates, automation explains about 14.5% of the total demand shifts, while new tasks explain about 2.1%, and sectoral TFP and markups explain 1.4% and 0.1%, respectively. Close to 82% of relative demand shifts remain unexplained. Since, as we have argued, factor-augmenting technologies are unlikely to contribute much to these between-group shifts, our framework highlights a puzzle: a sizable share of the implied relative demand shifts in the US economy since 1980 remains unexplained.

8.6 The incidence puzzle

Our reduced-form evidence revealed sizable effects of automation and new tasks on wages and employment. A natural way to think about employment effects is to introduce an endogenous labor supply margin so that demand shifts induce moves along an upward-sloping labor supply curve. For example, we may posit that the quantity of labor from skill group g is determined according to the labor supply schedule

$$\ell_g = \chi_g \cdot w_g^\varepsilon,$$

where $\varepsilon \geq 0$ is the net elasticity of labor supply (inclusive of income effects) and m_g a supply shifter. The case of inelastic labor supply studied so far is obtained when $\varepsilon = 0$. This labor supply curve can be the result of frictions (as in Kim and Vogel, 2021) or derived from household optimization with quasi-linear preferences (as in Acemoglu and Restrepo, 2022).

Proposition 9 extends to this environment, with now

$$d \ln w = \Theta^* \cdot \text{stack} \left(d \ln y + \sum_i \omega_{gi} \cdot z_{gi} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right),$$

where the propagation matrix, inclusive of endogenous supply responses, takes the form

$$\Theta^* = \frac{1}{\lambda + \varepsilon} \cdot \left(\mathbb{1} - \frac{1}{\lambda + \varepsilon} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \right)^{-1}.$$

The key difference with the previous matrix is that in place of λ , we have $\lambda + \varepsilon$. This extra term captures the intuitive fact that wage effects are less pronounced when labor supply is elastic since more of the adjustment takes place via quantities. Endogenous labor supply responses also weaken ripple effects, as lower hours worked for (negatively) affected groups means less competition for marginal tasks.

The incidence puzzle is that for realistic values of the labor supply elasticity, it is hard to make sense of the sizable reduced-form coefficients on our task variables. There are two ways of seeing the problem. First, as in a standard incidence analysis (and ignoring all general equilibrium interactions), the effect of a 1% decline in labor demand (measured as the shift in quantity demanded at constant prices) should be to reduce wages and employment by

$$\begin{aligned} d \ln w_g &= - \frac{1}{\sigma_g + \varepsilon} \cdot \text{shift in demand} \\ d \ln \ell_g &= - \frac{\varepsilon}{\sigma_g + \varepsilon} \cdot \text{shift in demand}, \end{aligned}$$

where σ_g is the demand elasticity for group g labor. This elasticity exceeds λ in our model, and so the incidence of a demand shock on wages and employment must be bounded above by $\frac{1}{\lambda + \varepsilon}$ and $\frac{\varepsilon}{\lambda + \varepsilon}$, respectively. Ripple effects and

other forces should, if anything, dampen the incidence of demand shocks, which means that these are upper bounds.

Alternatively, one can follow the derivations in Section 2, which imply that the row sum of Θ^* should be less than or equal to $1/(\lambda + \varepsilon)$. This places the same bound on our reduced-form estimates of the incidence of demand shocks on wages and employment.³⁸

The value of $\lambda = 0.5$ from Humlum (2020) and the estimate for $\varepsilon = 0.5$ reported in Chetty et al. (2011) yields an upper bound on the incidence rate of 1 for wages and 0.5 for employment, both of which are exceeded by our empirical estimates in Tables 1 and 2, centered around 1.25 for wages and 1.4 for hours worked per person. The root of the puzzle is the large estimates for employment. One could make sense of the estimated incidence on wages by positing lower values for λ or ε , but this would still predict an incidence in employment below 1.

The incidence problem is neither a technical problem nor an entirely new one. Rather, it reflects the fact that with elastic labor supply responses, it becomes impossible to generate large wage changes in general, as most of the adjustment is in quantities rather than prices.

We conjecture that both puzzles are related to the assumption that labor markets are fully competitive, and introducing non-competitive elements would provide at least a partial solution to both puzzles. For example, when the labor market is non-competitive, the implied relative demand shifts could be a significant exaggeration of the true changes in relative demand, which could be one reason why there appears to be a missing technology puzzle, and why employment responses are larger than predicted by the competitive benchmark. Relatedly, the presence of rents (wages that are above the opportunity cost of labor) for some groups, for instance as in Acemoglu and Restrepo (2024), would multiply the effects of automation on wages but also shift the economy off the labor supply curve. Such non-competitive elements could also amplify task displacement because they can induce additional automation as a means of dissipating rents accruing to certain worker groups.

9 Conclusion

This paper has reviewed and extended the recent literature on the task framework, where the production process is explicitly modeled as being based on the allocation of a range of tasks to different factors of production.

³⁸ The corresponding equation for employment becomes

$$d \ln \ell_g = \varepsilon \cdot \Theta^* \cdot \text{stack} \left(d \ln y + \sum_i \omega_{gi} \cdot z_{gi} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right).$$

The reduced-form estimates are now bounded above by ε times the row sums of Θ^* , which are less than $\frac{\varepsilon}{\lambda + \varepsilon}$.

The task model provides an attractive tool for studying ongoing labor market transformations in the United States and other industrialized nations for several reasons. To start with, an essential aspect of these transformations appears to be related to large changes in the nature of tasks—and occupations—that different types of workers perform in the labor market. Moreover, both the wage and occupational changes appear to be related to the rollout of new automation technologies that have substituted capital equipment and algorithms for tasks previously performed by some worker groups (Autor et al., 2003; Acemoglu and Autor, 2011; Acemoglu and Restrepo, 2022). Less appreciated but equally important are the effects of new technologies that have introduced new tasks for certain worker groups, ranging from new technical occupations to those based on digital tools, such as programming, design, integration functions and related service responsibilities (Lin, 2011; Acemoglu and Restrepo, 2018b; Autor et al., 2024). Automation and the introduction of new tasks cannot be easily studied in existing frameworks, which typically focus on factor-augmenting technological advances and do not distinguish the effects of different types of technologies.

The task framework not only adds descriptive realism to the modeling of the production process and the labor market, but leads to new comparative statics concerning the effects of technologies on the labor market. These new results are rooted in the *extensive-margin* effects of new technologies—that is, the reallocation of tasks away from certain worker groups as well as the reinstatement of some groups into new tasks—at given wages. We represent these extensive-margin influences via (direct) task displacement caused by automation and reinstatement generated by new tasks, and theoretically establish that they are very different than the consequences of technologies that make workers more productive in tasks they already perform or general factor-augmenting technologies that make factors uniformly more productive in all tasks.

The theoretical analysis in this chapter also builds a natural bridge between theory and empirics, and we expository and utilized this bridge at two different levels. The first is via a set of reduced-form equations that can be estimated to link relative wage (and employment) changes at the level of skill groups (e.g., groups distinguished by education, gender, age, ethnicity, etc.) to empirical measures of direct task displacement and reinstatement, as well as proxies for factor-augmenting technologies and sectoral reallocations. When estimated via reduced-form methods, this empirical framework points to a significant role of task displacement and reinstatement in accounting for the changes in the US wage and employment structure—in all cases explaining more than 50% of the variation between 1980 and 2016. In contrast, our proxies for other technological factors appear much less important in the distributional changes observed since 1980. This reduced-form evidence thus suggests that the extensive-margin effects of new technologies, typically ignored or bundled with other factors in standard approaches, should be the main focus when exploring the determinants of the recent evolution of the wage structure in the US and other industrialized economies.

Despite their simplicity and tight connection to theory, reduced-form equations have important limitations. First, they ignore the ripple effects that result from the spillovers from the technological changes impacting other worker groups. Second, reduced-form models are only informative about relative wage changes because productivity effects are subsumed into the constant term of the regression. Third, while the task displacement and reinstatement terms can be reasonably well approximated with the data we have available, our proxies for other technological influences may be less reliable. These shortcomings are rectified by a more structural approach that the task framework also enables—and we derived systematically from the multi-sector version of the model.

Specifically, the framework shows that the full effects of technological developments can be summarized by the following channels: a productivity effect, the direct extensive-margin effects on task allocations, task-price substitution effects, sectoral reallocations triggered by the uneven incidence of the technology in question across sectors, and the ripple effects. The ripple effects can be summarized (up to a first-order approximation) by a propagation matrix, which we develop and estimate via GMM from the same wage and task displacement and reinstements data. The remaining effects can be disciplined with external information on the elasticity of substitution between tasks within a sector and the elasticity of substitution between the outputs of different industries in the production of the final good.

Using this structural approach, our estimates of the propagation matrix and external estimates on the relevant elasticities, we carry out a full general equilibrium decomposition of the contribution of different technologies. We once again conclude that more than 50% of the changes in the US wage structure between 1980 and 2016 are driven by automation and new tasks.

One of the attractive features of the task framework is its flexibility, which we illustrated by showing how complex economic interactions can be modeled within this framework. There are several other directions for future work, which we hope our chapter will encourage:

- In this chapter, we focused on competitive models, with the exception of the exogenous sectoral markups that were introduced in the multi-sector model. The task framework naturally allows for the modeling of various imperfections. For example, the allocation of tasks to factors can be frictional due to search and matching considerations, discrimination against some groups in certain tasks or licensing. Additionally, the task model allows for efficiency-wage type considerations, rent-sharing, or explicit bargaining at the task level (e.g., Acemoglu and Restrepo, 2024). Such frictions not only cause inefficient assignment of tasks to factors, but also significantly enrich the effects of automation technologies, because these now have the additional role of dissipating worker rents and the adoption of these technologies can take place inefficiently as a result of employers' efforts to avoid paying worker rents. As mentioned above, non-competitive approaches can also hold the key to resolving the two puzzles we highlighted at the end of the previous section.

- More general preference structures, for example, including non-homothetic utility over different goods and services can be easily incorporated into this framework in order to study the process of structural change in the economy and its implications for the labor market. Such an extension can enable a more holistic analysis of the joint process of structural transformation and inequality following different types of technological influences.
- The task framework is ideally suited to studying the implications of trade in goods and services, offshoring and reshoring, and can be developed in the context of a multi-country setup in a relatively tractable form (Kikuchi, 2024).
- The task framework can be useful for exploring the effects of immigration and related changes on the supply side, making explicit how the effects of these developments depend on which tasks new or expanded labor groups compete for. For example, the framework suggests that the implications of an immigration shock should be very different when immigrants perform complementary tasks to natives; when they compete against machines; and when they compete for the tasks that certain native skill groups were previously performing.
- A major economic transformation will likely result from the rollout of new artificial intelligence (AI) tools in the coming decades. There is considerable uncertainty about the extent to which AI will be used to automate tasks, whether it can create new labor-intensive tasks and the magnitude of its productivity effects. It is also likely that developments in the AI industry can change product market competition and markups. These considerations increase the benefits of the task framework applied to study AI's variegated effects on the labor market (see, for example, Acemoglu, 2024; Acemoglu et al., 2022; Babina et al., 2024).
- The empirical work reported in this chapter uses publicly-available data, though we also mentioned an emerging literature using firm-level data. There is much more to be done with firm-level data and matched firm-worker data to investigate how task displacement and reinstatements take place and how this triggers a series of indirect effects, as not just the factors of production but also as firms compete with each other following the uneven adoption of various technologies.
- This chapter highlighted the importance of new tasks, which are challenging to measure in practice, and emphasized that future empirical work on the measurement of new tasks and their effects on different labor groups is an important direction (see Autor et al., 2024, for recent work on this).
- Finally, it would be useful to extend the theoretical and empirical approaches reviewed in this chapter, which relied on first-order approximations in order to incorporate the higher-order, nonlinear effects from large changes in technology or supplies.

Appendix A Equilibrium existence and uniqueness

This section proves Proposition 1, establishing the uniqueness of the equilibrium.

We first derive the equilibrium conditions in the text and provide a lemma for the Jacobian of task shares that will be used to establish the uniqueness of the equilibrium.

Preliminaries This section derives the equilibrium conditions E1-E5. E1 and E2 follow from cost minimization. For E3, note that the production of the final good is competitive, so task prices equal their marginal product $p(x) = M^{-1/\lambda}$. $\left(\frac{y}{y(x)}\right)^{1/\lambda}$, and

$$y(x) = \frac{1}{M} \cdot y \cdot p(x)^{-\lambda}. \quad (\text{A46})$$

For tasks in $\mathcal{T}_g(w)$, Eq. (A46) implies

$$\underbrace{A_g \cdot \psi_g(x) \cdot \ell_g(x)}_{y(x)} = \frac{1}{M} \cdot y \cdot \underbrace{\left(\frac{w_g}{A_g \psi_g(x)}\right)^{-\lambda}}_{p(x)},$$

which can be rearranged into E3. The same steps establish the corresponding equation for capital.

E4 imposes labor market clearing.

For E5, we multiply Eq. (A46) by p_x and integrate

$$\underbrace{\int y(x) \cdot p(x) \cdot dx}_y = \int_{\mathcal{T}} p_x \cdot y_x \cdot dx = \frac{1}{M} \cdot y \cdot \int_{\mathcal{T}} p_x^{1-\lambda} \cdot dx.$$

Canceling y on both sides yields the ideal-price index equation E5.

The Jacobian lemma The following lemma will be used in our proofs.

Lemma A1. *Let $\mathcal{H} = \mathbb{1} - \frac{1}{\lambda} \frac{\partial \ln \Gamma(w)}{\partial \ln w}$. For all wage vectors w , the matrix Σ is non-singular. Moreover, \mathcal{H} is a P-matrix of the Leontief type (i.e., with non-positive off-diagonal entries) whose inverse has all entries that are non-negative.*

Proof. Assumption 1 ensures that task shares are continuous and differentiable functions of wages. We now establish the properties of \mathcal{H} .

First, because $\partial \Gamma_g(w) / \partial w_{g'} \geq 0$ for $g' \neq g$, \mathcal{H} is a Z-matrix (it has negative off diagonals).

Second, \mathcal{H} has a positive dominant diagonal. This follows from the fact that $\mathcal{H}_{gg} = 1 - \frac{1}{\lambda} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_g} > 0$, and $\mathcal{H}_{gg} - \sum_{g' \neq g} |\mathcal{H}_{gg'}| = 1 - \sum_{g'} \frac{1}{\lambda} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}} >$

1. This last inequality follows because $\sum_{g'} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}} \leq 0$: when all wages rise by the same amount, workers lose tasks to capital but do not experience task reallocation among themselves.

Third, all eigenvalues of \mathcal{H} have real parts that exceed 1. This follows from Gershgorin's circle theorem: for each eigenvalue ζ of \mathcal{H} , we can find a dimension g such that $||\zeta - \mathcal{H}_{gg}|| < \sum_{g' \neq g} |\mathcal{H}_{gg'}|$. This inequality implies $\Re(\zeta) \in [\mathcal{H}_{gg} - \sum_{g' \neq g} |\mathcal{H}_{gg'}|, \mathcal{H}_{gg} + \sum_{g' \neq g} |\mathcal{H}_{gg'}|]$. Because $\mathcal{H}_{gg} - \sum_{g' \neq g} |\mathcal{H}_{gg'}| > 1$ for all g , all eigenvalues of \mathcal{H} have real parts greater than 1.

Fourth, since \mathcal{H} is a Z -matrix whose eigenvalues have positive real parts, it is also an M -matrix and a P -matrix of the Leontief type. The inverse of such matrices exists and has non-negative real entries. \square

Proof of Proposition 1. The derivations for the market-clearing wage in (4) were presented in the text.

The numeraire condition in (5) is obtained by substituting the expression for prices in E1 into the ideal price index in E5.

We now turn to existence and uniqueness. To prove that (4) and (5) admit a unique solution, we first show that, given a level for output y , there is a unique set of wages $\{w_g(y)\}_g$ that satisfies the market clearing conditions in (2). We then show there is a unique level of output that satisfies (5) evaluated at $\{w_g(y)\}_g$.

For the first step, Assumption 1 implies that $\Gamma_g(w)$ lies in a compact set $[\underline{\Gamma}, \bar{\Gamma}]$. $\mathbb{T} : w \rightarrow (\mathbb{T}w_1, \dots, \mathbb{T}w_G)'$ defined by $\mathbb{T}w_g = \left(\frac{y}{\ell_g}\right)^{\frac{1}{\lambda}} \cdot A_g^{1-1/\lambda} \cdot \Gamma_g(w)^{\frac{1}{\lambda}}$ for $g = 1, 2, \dots, G$ is a continuous mapping from the compact convex set $\mathbb{X} = \prod_{g=1}^G [(y/\ell_g)^{\frac{1}{\lambda}} \cdot A_g^{1-1/\lambda} \cdot \underline{\Gamma}^{\frac{1}{\lambda}}, (y/\ell_g)^{\frac{1}{\lambda}} \cdot A_g^{1-1/\lambda} \cdot \bar{\Gamma}^{\frac{1}{\lambda}}]$ onto itself. The existence of a positive wage vector $\{w_g(y)\}_g$ solving this fixed-point problem follows from Brouwer's fixed point theorem.

We now turn to uniqueness of $\{w_g(y)\}_g$. We can rewrite the system of equations $\{w_g(y)\}_g$ defining $\{w_g(y)\}_g$ in logs as $F(x) = \frac{1}{\lambda} \cdot \text{stack}(\ln y - \ln \ell_g)$, where $x = (\ln w_1, \dots, \ln w_G)$ and $F(x) = (f_1(x), \dots, f_G(x))$ with $f_g(x) = x_g - \frac{1}{\lambda} \cdot \ln \Gamma_g(x) - (1 - \frac{1}{\lambda}) \cdot \ln A_g$.

The Jacobian of F is given by the M -matrix \mathcal{H} . Theorem 5 from Gale and Nikaido (1965) shows that the solution to the system $F(x) = a$ is unique if the Jacobian of F is a P -matrix of the Leontief type. The theorem also shows that the unique solution $x(a)$ is increasing in a . As a result, the unique solution to the system of equations in (4) is $\{w_g(y)\}_g$ with $w_g(y)$ strictly increasing in y . We also note that $(y/\ell_g)^{1/\lambda} \cdot A_g^{1-1/\lambda} \cdot \underline{\Gamma}^{1/\lambda} \leq (y/\ell_g)^{1/\lambda} \cdot A_g^{1-1/\lambda} \cdot \bar{\Gamma}^{1/\lambda}$, so that $w_g(y) \rightarrow \infty$ as $y \rightarrow \infty$, and $w_g(y) \rightarrow 0$ as $y \rightarrow 0$.

To conclude, we show that there is a unique y that satisfies the ideal-price index equation (5). This condition can be written as $F(y) = 1$, where

$$F(y) = \left(\frac{1}{M} \int_{\mathcal{T}} \left[\min \left\{ \min_g \left\{ \frac{w_g(y)}{A_g \cdot \psi_g(x)} \right\}, \frac{1}{A_k \cdot \psi_k(x)} \right\} \right]^{1-\lambda} \cdot dx \right)^{1/(1-\lambda)}.$$

Because wages are increasing in y , $F(y)$ is also increasing in y . Assumption 1 also ensures that a positive mass of tasks must be allocated to labor at any wage level, which implies that $F(y)$ is increasing in y . The function $F(y)$ can be written as

$$F(y) = \left((A_k^{\lambda-1} \cdot \Gamma_k(w(y)) + \sum_g A_g^{\lambda-1} \cdot \Gamma_g(w(y)) \cdot w_g(y)^{1-\lambda}) \right)^{1/(1-\lambda)}.$$

As $y \rightarrow \infty$, $\Gamma_g(w) \cdot w_g(y)^{1-\lambda} \rightarrow \infty$ (since $\Gamma_g(w)$ is bounded from below and $\lambda < 1$) and $\Gamma_k(w(y)) \geq 0$. This implies $F(y) \rightarrow \infty$. Moreover, as $y \rightarrow 0$, $\Gamma_g(w) \cdot w_g(y)^{1-\lambda} \rightarrow 0$ (since $\lambda < 1$) and $\Gamma_k(w(y)) = 0$ (since, by Assumption 1, all tasks can be produced by at least one type of worker). This implies $F(y) \rightarrow 0$.

Because $F(y)$ is increasing in y , there is a unique $y \in (0, \infty)$ for which $F(y) = 1$ and, therefore, a unique equilibrium with wages $w_g = w_g(y)$. The equilibrium wages and the tie-breaking rule for tasks where there is indifference uniquely determine the task allocation.

Our argument for uniqueness also shows that, under Assumption 1, the unique equilibrium features finite output, positive wages, and positive task shares for all workers. Moreover, from $F(y) = 1$, we obtain that, in equilibrium, $1 - A_k^{\lambda-1} \cdot \Gamma_k(w) > 0$. \square

Appendix B Effects of technology

This section provides formulas for the effects of technology on wages.

Our comparative statics involve characterizing the change in task shares and equilibrium objects in response to infinitesimal changes in technology. For augmenting technologies this can be done via traditional differentiation, considering infinitesimal changes in $\psi_g(x)$, $\psi_k(x)$, A_g or A_k . Automation and new tasks creation, on the other hand, correspond to *discrete* shifts in capital and labor productivities over sets of positive or infinitesimal measure (e.g., capital becoming much more productive in many or a few tasks). In this Appendix, we define the notion of total derivatives of task shares with respect to these changes, which we use in the text. This definition applies to both the economy with and without ripples.

Let us write task shares in general as $\Gamma_g(\Psi)$, where Ψ designates all relevant parameters, including factor-augmenting terms, the A_g 's, and the measure of tasks M , with respect to which derivatives are defined in the usual manner. In the economy with ripples, one may also include wages as part of Ψ .

Consider a “small” (possibly *infinitesimal*) change in technology and wages. This change can be described as follows. Fix a small ϵ (so that infinitesimal changes correspond to $\epsilon \rightarrow 0$):

i. The automation of tasks in the set \mathcal{A}_g , with (Lebesgue) measure $\mathcal{O}(\epsilon)$ (i.e., there exists a constant \bar{c} such that the measure of \mathcal{A}_g is less than $\bar{c}\epsilon$). In this case, the quantity

$$a_g = \frac{1}{M} \int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} dx$$

gives the infinitesimal change in the task share of group g due to automation and

$$r_g = \frac{1}{M} \int_{\mathcal{A}_g} \psi_k^{auto}(x)^{\lambda-1} dx$$

gives the infinitesimal change in the task share of capital due to the automation of tasks in \mathcal{A}_g , where $\psi_k^{auto}(x)$ is the productivity of capital in task $x \in \mathcal{A}_g$ after the change in automation technology.

ii. The creation of new tasks in the set \mathcal{N}_g , with (Lebesgue) measure $\mathcal{O}(\epsilon)$. The quantity

$$n_g = \frac{1}{M} \int_{\mathcal{N}_g} \psi_g^{new}(x)^{\lambda-1} dx$$

gives the infinitesimal change in the task share of group g due to new tasks, with $\psi_g^{new}(x)$ being the productivity of labor of type g in tasks $x \in \mathcal{N}_g$ after the creation of new tasks.

iii. The change in Ψ , $d\Psi$, which is assumed to be of $\mathcal{O}(\epsilon)$ (i.e., there exists a constant \bar{c} such that $\|d\Psi\|$ is less than $\bar{c}\epsilon$).

Our notion of total derivatives of task shares is based on these quantities. In particular, define the total derivative of $\Gamma_g(\Psi)$ with respect to these infinitesimal changes as

$$d\Gamma_g(\Psi) = -a_g + n_g + \frac{\partial \Gamma_g}{\partial \Psi} \cdot d\Psi.$$

We show next that, just like the standard notion of total derivatives, this total derivative approximates the change in task shares with an error of order $o(\epsilon)$, meaning that it goes to zero faster than ϵ as ϵ goes to zero.

Likewise, define the total derivative of $\Gamma_k(\Psi)$ with respect to these infinitesimal changes as

$$d\Gamma_k(\Psi) = \sum_g r_g + \frac{\partial \Gamma_k}{\partial \Psi} \cdot d\Psi.$$

Moreover, the total derivative of any differentiable function $h(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi)$ can be determined via the chain rule as

$$\begin{aligned} dh(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi) &= \sum_g \frac{\partial h}{\partial \Gamma_g} \left(-a_g + n_g + \frac{\partial \Gamma_g}{\partial \Psi} \cdot d\Psi \right) \\ &\quad + \frac{\partial h}{\partial \Gamma_k} \left(\sum_g r_g + \frac{\partial \Gamma_k}{\partial \Psi} \cdot d\Psi \right) + \frac{\partial h}{\partial \Psi} \cdot d\Psi. \end{aligned}$$

The next lemma shows that, as for traditional derivatives, the total derivatives defined here for task shares—and via the chain rule for functions of task shares—provide a first-order approximation (in ϵ) to the change in $h(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi)$.

Lemma A2. *Let $h = h(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi)$. Suppose $h(\cdot)$ and $\Gamma(\cdot)$ are differentiable in Ψ (both before and after the change in technology). Suppose the sets \mathcal{A}_g and \mathcal{N}_g have Lebesgue measures $\mathcal{O}(\epsilon)$ and Ψ changes by $d\Psi$ of order $\mathcal{O}(\epsilon)$. Then the total change in h satisfies*

$$h' - h = \sum_g \frac{\partial h}{\partial \Gamma_g} \cdot \left(-a_g + \frac{\partial \Gamma_g}{\partial \Psi} \cdot d\Psi \right) + \frac{\partial h}{\partial \Gamma_k} \cdot \left(\sum_g r_g + \frac{\partial \Gamma_k}{\partial \Psi} \cdot d\Psi \right) + \frac{\partial h}{\partial \Psi} \cdot d\Psi + o(\epsilon). \quad (\text{A47})$$

Proof. We show this for automation. New tasks can be handled in the same manner. Note that

$$h_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi)_g, \Gamma_k^{\cup_g \mathcal{A}_g}(\Psi + d\Psi), \Psi + d\Psi - h(\{\Gamma_g(\Psi)\}, \Gamma_k(\Psi), \Psi),$$

where the notation $\Gamma_g^{\setminus \mathcal{A}_g}$ indicates that the task share is now computed over the set $\mathcal{T}_g \setminus \mathcal{A}_g$. The notation $\Gamma_k^{\cup_g \mathcal{A}_g}$ indicates that the task share of capital is now computed over the set $\mathcal{T}_k \cup_g \mathcal{A}_g$. This expression uses the fact that tasks in \mathcal{A}_g are automated in equilibrium by assumption.

A first-order Taylor expansion of h around $\{\Gamma_g\}$, Γ_k and Ψ yields

$$h' - h = \sum_g \frac{\partial h}{\partial \Gamma_g} \cdot \left(\Gamma_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_g(\Psi) \right) + \frac{\partial h}{\partial \Gamma_k} \cdot \left(\Gamma_k^{\cup_g \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_k(\Psi) \right) + \frac{\partial h}{\partial \Psi} \cdot d\Psi + o(\epsilon).$$

This step uses the fact that $\Gamma_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_g(\Psi)$ and $\Gamma_k^{\cup_g \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_k(\Psi)$ and $d\Psi$ are all $\mathcal{O}(\epsilon)$, so that the approximation error in the Taylor expansion is $o(\epsilon)$. This follows from $\Gamma_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi) = \Gamma_g^{\setminus \mathcal{A}_g}(\Psi) + \mathcal{O}(\epsilon)$ (from continuity), which implies $\Gamma_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_g(\Psi) = \Gamma_g^{\setminus \mathcal{A}_g}(\Psi) - \Gamma_g(\Psi) + \mathcal{O}(\epsilon)$. The right side is $\mathcal{O}(\epsilon)$ because $\Gamma_g^{\setminus \mathcal{A}_g}(\Psi) - \Gamma_g(\Psi)$ differ over a set of measure $\mathcal{O}(\epsilon)$. The argument for $\Gamma_k^{\cup_g \mathcal{A}_g}(\Psi + d\Psi) - \Gamma_k(\Psi)$ is the same.

A second first-order Taylor expansion, this time of the task shares $\Gamma_g^{\setminus \mathcal{A}_g}(\Psi + d\Psi)$ and $\Gamma_k^{\cup_g \mathcal{A}_g}(\Psi + d\Psi)$ around Ψ gives

$$h' - h = \sum_g \frac{\partial h}{\partial \Gamma_g} \cdot \left(-a_g + \frac{\partial \Gamma_g^{\setminus \mathcal{A}_g}}{\partial \Psi} \cdot d\Psi \right) + \frac{\partial h}{\partial \Gamma_k} \cdot \left(\sum_g r_g + \frac{\partial \Gamma_k^{\cup_g \mathcal{A}_g}}{\partial \Psi} \cdot d\Psi \right) + \frac{\partial h}{\partial \Psi} \cdot d\Psi + o(\epsilon). \quad (\text{A48})$$

We now show that $\frac{\partial \Gamma_g^{\setminus \mathcal{A}_g}}{\partial \Psi} \cdot d\Psi = \frac{\partial \Gamma_g}{\partial \Psi} \cdot d\Psi + o(\epsilon)$ and $\frac{\partial \Gamma_k^{\cup_g \mathcal{A}_g}}{\partial \Psi} \cdot d\Psi = \frac{\partial \Gamma_k}{\partial \Psi} \cdot d\Psi + o(\epsilon)$. We establish this claim by considering the different elements in Ψ one by one. Changes in wages and uniformly augmenting technologies only affect task shares by reallocating marginal tasks. By assumption \mathcal{A}_g is in the interior of \mathcal{T}_g and all of the tasks in this set are strictly cheaper when automated (i.e., are not marginal). This implies $\frac{\partial \Gamma_g^{\setminus \mathcal{A}_g}}{\partial w} = \frac{\partial \Gamma_g}{\partial w}$ and $\frac{\partial \Gamma_k^{\cup_g \mathcal{A}_g}}{\partial w} = \frac{\partial \Gamma_k}{\partial w}$. The same logic implies that for factor-augmenting technologies A_g , we have $\frac{\partial \Gamma_g^{\setminus \mathcal{A}_g}}{\partial A_{g'}^{\setminus \mathcal{A}_g}} = \frac{\partial \Gamma_g}{\partial A_{g'}^{\setminus \mathcal{A}_g}}$ and $\frac{\partial \Gamma_k^{\cup_g \mathcal{A}_g}}{\partial A_{g'}^{\setminus \mathcal{A}_g}} = \frac{\partial \Gamma_k}{\partial A_{g'}^{\setminus \mathcal{A}_g}}$. For augmenting technologies at the intensive margin, the set of automated tasks and the set of tasks with productivity improvements may overlap. However, the improvements are $\mathcal{O}(\epsilon)$ and the range of overlap is $\mathcal{O}(\epsilon)$, which means that the overlap is $\mathcal{O}(\epsilon^2)$, which is at least as fast as $o(\epsilon)$ as claimed. Substituting these back into (A48) gives (A47). \square

Remark 1. The proof uses the fact that all tasks in \mathcal{A}_g become automated. The assumption that $\pi_g^{\text{auto}}(x) > 0$ ensures this, because, at the initial equilibrium wages, producing these tasks with capital is cheaper than assigning them to labor. Because the change in wages is also small, the same remains true in the new equilibrium. Note that this logic can fail for large automation shocks, in which case only a subset of tasks in \mathcal{A}_g may become automated in equilibrium.

Remark 2. Applying the lemma to $h = \Gamma_g(\Psi)$ or $h = \Gamma_k(\Psi)$ shows that our definition of derivatives provides a first-order approximation to the change in $\Gamma_g(\Psi)$ and $\Gamma_k(\Psi)$ whose error term is $o(\epsilon)$.

Remark 3. If h , Γ_g , and Γ_k are twice differentiable, then the same steps establish the sharper bound

$$h' - h = dh(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi) + \mathcal{O}(\epsilon^2).$$

This means that the derivative $dh(\{\Gamma_g(\Psi)\}_g, \Gamma_k(\Psi), \Psi)$ approximates the change in $h' - h$ with a small approximation error that goes to zero no slower than ϵ^2 .

Remark 4. Equilibrium wages are one of the variables in Ψ and our expressions so far assume that changes in wages are also $\mathcal{O}(\epsilon)$. We show here that this is indeed the case. In particular, note that equilibrium wages solve a system of the form

$$h(\{\Gamma_g(\Psi_0, \Omega)\}_g, \Gamma_k(\Psi_0, \Omega), \Psi_0, \Omega) = 0, \quad (\text{A49})$$

where, for emphasis, we have separated wages from exogenous technological parameters in Ψ_0 . Recall also that h is differentiable in w (by virtue of Assumption 1) and the Jacobian of h with respect to w , denoted by J_w is non-singular—a consequence of the uniqueness of the equilibrium established in Proposition 1.

Consider a change in automation, new tasks, and other technologies of order ϵ , and denote the new equilibrium wage by $w' = w + dw$. Applying Lemma A2 to differentiate h with respect to automation, new tasks, and Ψ_0 (all changes of order $\mathcal{O}(\epsilon)$) yields

$$\begin{aligned} 0 = & \sum_g \frac{\partial h}{\partial \Gamma_g} \left(-a_g + n_g + \frac{\partial \Gamma_g}{\partial \Psi_0} \cdot d\Psi_0 \right) + \frac{\partial h}{\partial \Gamma_k} \left(\sum_g r_g + \frac{\partial \Gamma_k}{\partial \Psi_0} \cdot d\Psi_0 \right) \\ & + \frac{\partial h}{\partial \Psi_0} \cdot d\Psi_0 + R_1 + h(\{\Gamma_g(\Psi_0, \Omega + d\Omega)\}_g, \Gamma_k(\Psi_0, \Omega + d\Omega), \Psi_0, \Omega + d\Omega) \\ & - h(\{\Gamma_g(\Psi_0, \Omega)\}_g, \Gamma_k(\Psi_0, \Omega), \Psi_0, \Omega), \end{aligned} \quad (\text{A50})$$

where the approximation error R_1 is $o(\epsilon)$ and the derivatives in the first line are evaluated at the new equilibrium wages $w' = w + dw$. Taking limits in (A50) as $\epsilon \rightarrow 0$ implies

$$\begin{aligned} 0 = & h(\{\Gamma_g(\Psi_0, w + dw)\}_g, \Gamma_k(\Psi_0, w + dw), \Psi_0, w + dw) \\ & - h(\{\Gamma_g(\Psi_0, w)\}_g, \Gamma_k(\Psi_0, w), \Psi_0, w). \end{aligned}$$

By the continuity of h as a function of w , this equality can only hold if $dw \rightarrow 0$ as $\epsilon \rightarrow 0$. We finally show that $dw \rightarrow 0$ at the same rate as $\epsilon \rightarrow 0$, establishing the claim that dw is $\mathcal{O}(\epsilon)$. Suppose by way of contradiction that $\epsilon / \|dw\| \rightarrow 0$ as $\epsilon \rightarrow 0$, so that dw is not $\mathcal{O}(\epsilon)$. A Taylor expansion of the second line in (A50) around wages of order dw yields

$$\begin{aligned} 0 = & \sum_g \frac{\partial h}{\partial \Gamma_g} \left(-a_g + n_g + \frac{\partial \Gamma_g}{\partial \Psi_0} \cdot d\Psi_0 \right) + \frac{\partial h}{\partial \Gamma_k} \left(\sum_g r_g + \frac{\partial \Gamma_k}{\partial \Psi_0} \cdot d\Psi_0 \right) \\ & + \frac{\partial h}{\partial \Psi_0} \cdot d\Psi_0 + J_w \cdot dw + R_1 + R_2, \end{aligned}$$

where R_2 is $o(\|dw\|)$. Dividing both sides by $\|dw\|$ and taking limits as $\epsilon \rightarrow 0$ yields

$$\begin{aligned} 0 = & \sum_g \frac{\partial h}{\partial \Gamma_g} \left(-\frac{a_g}{\epsilon} \frac{\epsilon}{\|dw\|} + \frac{n_g}{\epsilon} \frac{\epsilon}{\|dw\|} + \frac{\partial \Gamma_g}{\partial \Psi_0} \cdot \frac{d\Psi_0}{\epsilon} \frac{\epsilon}{\|dw\|} \right) \\ & + \frac{\partial h}{\partial \Gamma_k} \left(\frac{\sum_g r_g}{\epsilon} \frac{\epsilon}{\|dw\|} + \frac{\partial \Gamma_k}{\partial \Psi_0} \cdot \frac{d\Psi_0}{\epsilon} \frac{\epsilon}{\|dw\|} \right) \\ & + \frac{\partial h}{\partial \Psi_0} \cdot \frac{d\Psi_0}{\epsilon} \frac{\epsilon}{\|dw\|} + J_w \cdot \frac{dw}{\|dw\|} + \frac{R_1}{\epsilon} \frac{\epsilon}{\|dw\|} + \frac{R_2}{\|dw\|} \end{aligned}$$

$$\Leftrightarrow 0 = \lim_{\epsilon \rightarrow 0} J_w \cdot \frac{dw}{||dw||}.$$

This is because the $||dw||$ in the denominators dominates all terms except $J_w \cdot dw$. Because J_w is non-singular, this yields a contradiction and we conclude that dw is of order $\mathcal{O}(\epsilon)$ as claimed.

We will use the lemma repeatedly in the appendix. In particular, our strategy is to totally differentiate equilibrium conditions to obtain a linear system in dw and dy (the change in wages and output) relating these to the infinitesimal changes in technology (summarized by a_g , n_g , r_g , and $d\Psi_0$). Lemma A2 implies that solving for dw and dy in this linear system approximates the equilibrium change with an error of order $o(\epsilon)$. The same lemma can be applied to the multi-sector economy, and there we also obtain linear equations for dw , dp , and dy (the change in wages, sectoral prices, and output) which can also be solved and provide a first order approximation to the equilibrium change in these endogenous objects.

B.1 No-ripple economy

This section derives the formulas for the effects of technology in the no-ripple economy in Propositions 2, 3, 4, and 5. We also provide formulas for the effects of these technologies on the labor share and output.

Proof of Proposition 2. Consider a new technology that automates tasks in \mathcal{A}_g .

To derive Eq. (9), we start from (4) and compute its total derivative

$$d \ln w_g = \frac{1}{\lambda} \cdot d \ln y - \frac{1}{\lambda} \frac{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx} = \frac{1}{\lambda} \cdot \left(d \ln y - d \ln \Gamma_g^{auto} \right).$$

To derive Eq. (10), we start from the definition of the cost function on the right-side of (5) (in logs). In equilibrium, $\ln C(w) = 0$. Computing its total derivative yields

$$\begin{aligned} d \ln C(w) &= \sum_g s_g^y \cdot d \ln w_g \\ &+ \sum_g \frac{1}{1-\lambda} \cdot \frac{1}{M} \cdot \left[\frac{s_K^y}{\Gamma_k} \cdot \int_{\mathcal{A}_g} \psi_k^{auto}(x)^{\lambda-1} \cdot dx - \frac{s_g^y}{\Gamma_g} \cdot \int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx \right]. \end{aligned}$$

The first term gives the effect of wage changes on cost, which is derived from Shephard's lemma.

Using the fact that $s_K^y = \Gamma_k \cdot A_k^{\lambda-1}$ and $s_g^y = \Gamma_g \cdot A_g^{\lambda-1} \cdot w_g^{1-\lambda}$, the change in costs can be rewritten as

$$\begin{aligned}
d \ln \mathcal{C}(w) &= \sum_g s_g^y \cdot d \ln w_g \\
&+ \sum_g \frac{1}{1-\lambda} \cdot \frac{1}{M} \cdot \left[\int_{\mathcal{A}_g} A_k^{\lambda-1} \cdot \psi_k^{\text{auto}}(x)^{\lambda-1} \cdot dx \right. \\
&\quad \left. - \int_{\mathcal{A}_g} A_g^{\lambda-1} \cdot \psi_g(x)^{\lambda-1} \cdot w_g^{1-\lambda} \cdot dx \right] \\
&= \sum_g s_g^y \cdot d \ln w_g - \sum_g A_g^{\lambda-1} \cdot w_g^{1-\lambda} \cdot \frac{1}{M} \int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot \pi_g^{\text{auto}}(x) \cdot dx \\
&= \sum_g s_g^y \cdot d \ln w_g \\
&- \sum_g \underbrace{\Gamma_g \cdot A_g^{\lambda-1} \cdot w_g^{1-\lambda}}_{s_g^y} \cdot \underbrace{\frac{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx}}_{d \ln \Gamma_g^{\text{auto}}} \cdot \underbrace{\frac{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot \pi_g^{\text{auto}}(x) \cdot dx}{\int_{\mathcal{A}_g} \psi_g(x)^{\lambda-1} \cdot dx}}_{\pi_g^{\text{auto}}},
\end{aligned}$$

which shows that $d \ln \mathcal{C}(w) = \sum_g s_g^y \cdot d \ln w_g - \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{auto}} \cdot \pi_g^{\text{auto}}$. In equilibrium, $d \ln \mathcal{C}(w) = 0$, which establishes (10).

We now provide expressions for output and the labor share. Solving for output from (9) and (10), we obtain

$$d \ln y = \sum_g \frac{s_g^y}{s_L^y} \cdot d \ln \Gamma_g^{\text{auto}} \cdot (1 + \lambda \cdot \pi_g^{\text{auto}}).$$

The change in the labor share can then be computed from $d \ln s_L^y = \frac{1}{s_L^y} \sum_g s_g^y \cdot d \ln w_g - d \ln y$ as

$$d \ln s_L^y = - \sum_g \frac{s_g^y}{s_L^y} \cdot (1 - (1 - \lambda) \cdot \pi_g^{\text{auto}}) \cdot d \ln \Gamma_g^{\text{auto}}.$$

Finally, the capital share can be obtained from $d \ln s_K^y = \frac{-ds_L^y}{s_K^y} = - \frac{s_L^y}{s_K^y} \cdot d \ln s_L^y$ as

$$d \ln s_K^y = \sum_g \frac{s_g^y}{s_K^y} \cdot (1 - (1 - \lambda) \cdot \pi_g^{\text{auto}}) \cdot d \ln \Gamma_g^{\text{auto}}. \quad \square$$

Proof of Proposition 3. To derive Eq. (12), we start from (4) and totally differentiate it to obtain

$$d \ln w_g = \frac{1}{\lambda} \cdot (d \ln y + d \ln \Gamma_g^{\text{new}} - d \ln M)$$

$$= \frac{1}{\lambda} \cdot \left(d \ln y + \frac{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx} - d \ln M \right).$$

To derive Eq. (13), we start from the definition of the cost function on the right-side of (5). As before, the change in log cost is

$$\begin{aligned} d \ln \mathcal{C}(w) &= \sum_g s_g^y \cdot d \ln w_g \\ &+ \sum_g \frac{1}{1-\lambda} \cdot \frac{1}{M} \left[\frac{s_g^y}{\Gamma_g} \cdot \int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx - \int_{\mathcal{N}_g} dx \right], \end{aligned}$$

where we used the fact that $d \ln M = \frac{1}{M} \sum_g \int_{\mathcal{N}_g} dx$. The first term gives the effect of wage changes on cost, which is derived from Shephard's lemma.

Using the fact that $s_g^y = \Gamma_g \cdot A_g^{\lambda-1} \cdot w_g^{1-\lambda}$, the change in costs can be rewritten as

$$\begin{aligned} d \ln \mathcal{C}(w) &= \sum_g s_g^y \cdot d \ln w_g \\ &+ \sum_g \frac{1}{1-\lambda} \cdot \frac{1}{M} \cdot \left[\int_{\mathcal{N}_g} A_g^{\lambda-1} \cdot \psi_g(x)^{\lambda-1} \cdot w_g^{1-\lambda} \cdot dx - \int_{\mathcal{N}_g} dx \right] \\ &= \sum_g s_g^y \cdot d \ln w_g - \sum_g \frac{1}{M} \cdot A_g^{\lambda-1} \cdot w_g^{1-\lambda} \cdot \int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot \pi^{\text{new}}(x) \cdot dx \\ &= \sum_g s_g^y \cdot d \ln w_g \\ &- \sum_g \underbrace{\Gamma_g \cdot A_g^{\lambda-1} \cdot w_g^{1-\lambda}}_{s_g^y} \cdot \underbrace{\frac{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx}}_{d \ln \Gamma_g^{\text{new}}} \cdot \underbrace{\frac{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot \pi^{\text{new}}(x) \cdot dx}{\int_{\mathcal{N}_g} \psi_g(x)^{\lambda-1} \cdot dx}}_{\pi_g^{\text{new}}} \end{aligned}$$

which shows that $d \ln \mathcal{C}(w) = \sum_g s_g^y \cdot d \ln w_g - \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{new}} \cdot \pi_g^{\text{new}}$. In equilibrium, $d \ln \mathcal{C}(w) = 0$, which establishes (13).

We now provide expressions for output and the labor share. Solving for output from (12) and (13), we obtain

$$d \ln y = \sum_g s_g^y \cdot d \ln \Gamma_g^{\text{new}} \cdot \left[1 - \frac{1}{s_L^y} + \left((1-\lambda) + \frac{1}{s_L^y} \cdot \lambda \right) \cdot \pi_g^{\text{new}} \right].$$

The change in the labor share can then be computed from $d \ln s_L^y = \frac{1}{s_L^y} \sum_g s_g^y \cdot d \ln w_g - d \ln y$ as

$$d \ln s_L^y = \frac{s_k^y}{s_L^y} \cdot \sum_g s_g^y \cdot d \ln \Gamma_g^{new} \cdot (1 + (1 - \lambda) \cdot \pi_g^{new})$$

Finally, the capital share can be computed from $d \ln s_K^y = \frac{-ds_L^y}{s_K^y} = -\frac{s_L^y}{s_K^y} \cdot d \ln s_L^y$ as

$$d \ln s_K^y = -\sum_g s_g^y \cdot d \ln \Gamma_g^{new} \cdot (1 + (1 - \lambda) \cdot \pi_g^{new}). \quad \square$$

Proof of Proposition 4. Differentiating Eq. (4) establishes (14):

$$\begin{aligned} d \ln w_g &= \frac{1}{\lambda} \cdot d \ln y + (1 - 1/\lambda) \cdot d \ln A_g \\ &+ (1 - 1/\lambda) \cdot \underbrace{\frac{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot d \ln \psi_g(x) \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx}}_{d \ln \psi_g^{\text{intensive}}}. \end{aligned}$$

Total differentiation of the cost function $\mathcal{C}(w)$ on the right-hand side of (5) implies

$$\begin{aligned} d \ln \mathcal{C}(w) &= \sum_g s_g^y \cdot d \ln w_g - \sum_g s_g^y \cdot d \ln A_g \\ &- \sum_g s_g^y \cdot \underbrace{\frac{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot d \ln \psi_g(x) \cdot dx}{\int_{\mathcal{T}_g^*} \psi_g(x)^{\lambda-1} \cdot dx}}_{d \ln \psi_g^{\text{intensive}}}, \end{aligned}$$

establishing (15). As before, the first term gives the effect of wage changes on cost, which is derived from Shephard's lemma.

We now provide expressions for output and the labor share. Solving for output from (14) and (15), we obtain

$$d \ln y = \sum_g \frac{s_g^y}{s_L^y} \cdot (d \ln A_g + d \ln \psi_g^{\text{intensive}}).$$

In this case, the labor share (and hence the capital share) remains unchanged. This follows from the fact that, in the no-ripple economy, \mathcal{T}_k does not change in response to labor-augmenting technologies. \square

Proof of Proposition 5. Total differentiation of Eq. (4) implies

$$d \ln w_g = \frac{1}{\lambda} \cdot d \ln y,$$

establishing (16).

Total differentiation of the cost function $\mathcal{C}(w)$ in the right-hand side of (5) implies

$$d \ln \mathcal{C}(w) = \sum_g s_g^y \cdot d \ln w_g - s_K^y \cdot d \ln A_k - s_K^y \cdot \underbrace{\frac{\int_{\mathcal{T}_k^*} \psi_k(x)^{\lambda-1} \cdot d \ln \psi_k(x) \cdot dx}{\int_{\mathcal{T}_k^*} \psi_k(x)^{\lambda-1} \cdot dx}}_{d \ln \psi_k^{\text{intensive}}},$$

establishing (17). As before, the first term gives the effect of wage changes on cost, which is derived from Shephard's lemma.

We now provide expressions for output and the labor share. Solving for output from (16) and (17), we obtain

$$d \ln y = \lambda \cdot \frac{s_L^y}{s_L^y} \cdot (d \ln A_k + d \ln \psi_k^{\text{intensive}}).$$

The change in the labor share can then be computed from $d \ln s_L^y = \frac{1}{s_L^y} \sum_g s_g^y \cdot d \ln w_g - d \ln y$ as

$$d \ln s_L^y = (1 - \lambda) \cdot \frac{s_L^y}{s_L^y} \cdot (d \ln A_k + d \ln \psi_k^{\text{intensive}}).$$

Finally, the capital share can be computed from $d \ln s_K^y = \frac{-ds_L^y}{s_K^y} = -\frac{s_L^y}{s_K^y} \cdot d \ln s_L^y$ as

$$d \ln s_K^y = -(1 - \lambda) \cdot (d \ln A_k + d \ln \psi_k^{\text{intensive}}).$$

□

B.2 Effects of technology with ripples

This section proves Proposition 6 and explains the details of how we apply it to characterize the effects of the different technologies. We then prove Proposition 7.

Proof of Proposition 6. Lemma A2 shows that we can totally differentiate (4) in response to an infinitesimal change in technology (or automation and new tasks in sets of infinitesimal measure) to obtain (18) in the main text, where z_g depends on the shocks considered. Stacking (18) and solving for wages gives (19).

Eq. (20) follows from the fact that $d \ln \mathcal{C}(w) = \sum_g s_g^y \cdot d \ln w - \pi$. As before, the first term gives the effect of wage changes on cost, which is derived from

Shephard's lemma. Note that here, π is computed as in the no-ripple economy, since it is by definition equal to the effect of technology holding wages constant.

The calculation of the effects of uniform-augmenting technologies in terms of the propagation matrix requires some further explanation. For uniform-labor augmenting improvements, differentiating (4) yields

$$\begin{aligned} d \ln w_g &= \frac{1}{\lambda} \cdot d \ln y + \underbrace{(1 - 1/\lambda) \cdot d \ln A_g - \frac{1}{\lambda} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \cdot d \ln A}_{z_g} \\ &\quad + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \cdot d \ln w, \end{aligned}$$

where $d \ln A = (d \ln A_1, \dots, d \ln A_G)$ and we used the fact that an increase in A_g generates an equal task reassignment as a commensurate decrease in w_g . Solving for $d \ln w$ yields $d \ln w = \Theta \cdot d \ln y + (1 - \Theta) \cdot d \ln A$, which is equivalent to the formula used in the text.

For uniform-capital augmenting improvements, differentiating (4) yields

$$d \ln w_g = \frac{1}{\lambda} \cdot d \ln y + \underbrace{\sum_{g'} \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}} \cdot d \ln A_k + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \cdot d \ln w}_{z_g}$$

where this expression uses the fact that an increase in A_k generates the same reallocation of tasks as an increase in all wages of the same magnitude. Solving for $d \ln w$ yields $d \ln w = \Theta \cdot (d \ln y + \lambda \cdot d \ln A_k) - d \ln A_k$, or equivalently $d \ln w_g = \rho_g \cdot d \ln y - (1 - \rho_g \cdot \lambda) \cdot d \ln A_k$ as claimed in the text. \square

Proof of Proposition 7. The expression for the change in wages in (24) follows from differentiating Eq. (4):

$$d \ln w_g = \frac{1}{\lambda} \cdot d \ln y - \frac{1}{\lambda} \cdot d \ln \ell_g + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w} \cdot d \ln w.$$

Stacking across groups and solving for $d \ln w_g$ yields (24).

The fact that there are no average wage changes follows from differentiating the cost function in (5). Because technology does not change, we have

$$d \ln \mathcal{C}(w) = \sum_g s_g^y \cdot d \ln w_g = 0,$$

which follows from Shephard's lemma. \square

Appendix C Equilibrium in the multi-sector economy

This section provides details and proofs for the multi-sector economy.

Preliminaries we first derive the equilibrium conditions E1-E6.

E1 and E2 follow from cost minimization.

For E3, because producers in sector i face an exogenous markup μ_i , they use task $x \in \mathcal{T}_i$ until $p_i \cdot M_i^{-1/\lambda} \cdot A_i^{1-1/\lambda} \cdot \left(\frac{y_i}{p(x)}\right)^{1/\lambda} = \mu_i \cdot p(x)$, so that the value of the task marginal product (on the left) exceeds its marginal cost (on the right) by a factor of μ_i . The quantity of task $x \in \mathcal{T}_i$ used is then

$$y(x) = y_i \cdot p_i^\lambda \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \frac{1}{M_i} \cdot p(x)^{-\lambda}. \quad (\text{A51})$$

For tasks in $\mathcal{T}_{gi}(w)$, Eq. (A51) implies

$$\underbrace{A_g \cdot \psi_g(x) \cdot \ell_g(x)}_{y(x)} = y_i \cdot p_i^\lambda \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \frac{1}{M_i} \cdot \underbrace{\left(\frac{w_g}{A_g \psi_g(x)}\right)^{-\lambda}}_{p(x)},$$

which explains E3. The same steps establish the corresponding equation for capital.

E4 imposes labor market clearing, now adding labor demand across all sectors.

For E5, multiply Eq. (A51) by $\mu_i \cdot p_x$ and integrate

$$\underbrace{\mu_i \cdot \int y(x) \cdot p(x) \cdot dx}_{y_i \cdot p_i} = \int_{\mathcal{T}_i} y_i \cdot p_i^\lambda \cdot \mu_i^{1-\lambda} \cdot A_i^{\lambda-1} \cdot \frac{1}{M_i} \cdot p(x)^{1-\lambda} \cdot dx.$$

Canceling y_i on both sides and solving for p_i gives the price index equation E5.

Finally, E6 follows the numeraire condition and requires the price of the final good to be 1.

Proofs for multi-sector model propositions We now prove Proposition 8 describing the equilibrium in the multi-sector economy and then turn to Propositions 9 and 10 characterizing the impact of technology and markups, respectively.

Proof of Proposition 8. We first derive the expression for the market-clearing wage in Eq. (25). Aggregating E3 across all tasks assigned to group g in all sectors, and using the definition of $\Gamma_{gi}(w)$, we can write the labor market clearing condition as

$$y \cdot A_g^{\lambda-1} \cdot w_g^{-\lambda} \cdot \left[\sum_i s_i^y(p) \cdot p_i^{\lambda-1} \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \Gamma_{gi}(w) \right] = \ell_g.$$

Isolating w_g from this equation yields (25).

The formula for sectoral prices in terms of task shares in (26) is obtained by substituting the expression for prices in E1 into the price index formula in E5.

The final equilibrium equation in (27) is just E6. \square

Proof of Proposition 9. Lemma A2 implies that we can totally differentiate (25) as

$$\begin{aligned} d \ln w_g &= \frac{1}{\lambda} \cdot d \ln y + \frac{1}{\lambda} \cdot \sum_i \omega_{gi} \cdot z_{gi} + \frac{1}{\lambda} \cdot (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \\ &\quad + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \cdot d \ln w. \end{aligned} \quad (\text{A52})$$

Stacking (A52) and solving for wages gives (28).

Eq. (29) follows from the fact that $d \ln p_i = d \ln \mathcal{C}_i(w) = \sum_g s_g^{y_i} \cdot d \ln w - \pi$, again from Shephard's lemma. Finally, Eq. (30) follows from the fact that $0 = d \ln c^f(p) = \sum_i s_i \cdot d \ln p_i$, again from Shephard's lemma, but applied to the production of the final good. \square

Proof of Proposition 10. Totally differentiating (25), we obtain

$$\begin{aligned} d \ln w_g &= \frac{1}{\lambda} \cdot d \ln y - \sum_i \omega_{gi} \cdot d \ln \mu_i + \frac{1}{\lambda} \cdot (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \\ &\quad + \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \cdot d \ln w. \end{aligned}$$

Stacking these equations for all groups and solving for wages gives (33).

Eq. (34) follows from the fact that $d \ln p_i = d \ln \mathcal{C}_i(w) = \sum_g s_g^{y_i} \cdot d \ln w + d \ln \mu_i$, again from Shephard's lemma.

Finally, Eq. (35) follows from the fact that $0 = d \ln c^f(p) = \sum_i s_i \cdot d \ln p_i$, again from Shephard's lemma, but applied to the production of the final good. \square

Appendix D Endogenous labor supply

The following proposition extends our analysis to a multi-sector economy with endogenous labor supply. For this proposition, we assume labor supply is given by $\ell_g = \chi_g \cdot w_w^\varepsilon$.

Proposition A1 (Effects of technology in the multi-sector economy). *With an endogenous labor supply, equilibrium wages w , industry prices p , and the level of output y , solve the system of equations*

$$w_g = \left(\frac{y}{\chi_g} \right)^{1/(\lambda+\varepsilon)} \cdot A_g^{(\lambda-1)/(\lambda+\varepsilon)} \cdot \left[\sum_i s_i^y(p) \cdot p_i^{\lambda-1} \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \Gamma_{gi}(w) \right]^{1/(\lambda+\varepsilon)} \quad \text{for } g \in \mathbb{G}, \quad (\text{A53})$$

$$p_i = \mu_i \cdot \underbrace{\frac{1}{A_i} \cdot \left(\Gamma_{ki}(w) \cdot A_k^{\lambda-1} + \sum_g \Gamma_{gi}(w) \cdot \left(\frac{w_g}{A_g} \right)^{1-\lambda} \right)^{1/(1-\lambda)}}_{\equiv \mathcal{C}_i(w)} \quad \text{for } i \in \mathbb{I}, \quad (\text{A54})$$

$$1 = c_f(p), \quad (\text{A55})$$

where $\mathcal{C}_i(w)$ denotes the marginal cost of producing output of sector i .

In addition, the effect of a change in technology with direct effect $\{z_{gi}\}_{g \in \mathbb{G}, i \in \mathbb{I}}$ and productivity gains $\{\pi_{gi}\}_{g \in \mathbb{G}, i \in \mathbb{I}}$ on wages, sectoral prices, and output is given by the formulas in Proposition 9, with the propagation matrix redefined as

$$\Theta^* = \frac{1}{\lambda + \varepsilon} \cdot \left(1 - \frac{1}{\lambda + \varepsilon} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \right)^{-1},$$

and direct effect re-scaled by $\lambda + \varepsilon$ (so that direct effect are $(1/(\lambda + \varepsilon)) \cdot z_{gi}$).

Proof. The equilibrium conditions in this case are still given by E1–E6. The only difference is that the market clearing condition in E4 is now

$$\sum_i \int_{\mathcal{T}_{gi}} \ell_g(x) \cdot dx = \chi_g \cdot w_g^\varepsilon.$$

Following the same steps as in the proof of Proposition 8, we can write this condition as

$$y \cdot A_g^{\lambda-1} \cdot w_g^{-\lambda} \cdot \left[\sum_i s_i^y(p) \cdot p_i^{\lambda-1} \cdot \mu_i^{-\lambda} \cdot A_i^{\lambda-1} \cdot \Gamma_{gi}(w) \right] \cdot = \chi_g \cdot w_g^\varepsilon.$$

Isolating w_g from this equation yields (A53).

The formula for sectoral prices in terms of task shares in (A54) is obtained by substituting the expression for prices in E1 into the price index formula in E5.

The final equilibrium equation in (A55) is E6.

We now show that the formulas for the effects of technology coincide with those in Proposition 9 with Θ^* in place of Θ .

Totally differentiating (A53) yields

$$d \ln w_g = \frac{1}{\lambda + \varepsilon} d \ln y + \frac{1}{\lambda + \varepsilon} \sum_i \omega_{gi} \cdot z_{gi} + \frac{(\lambda - \eta)}{\lambda + \varepsilon} \cdot \sum_i \omega_{gi} \cdot d \ln p_i$$

$$+ \frac{1}{\lambda + \varepsilon} \cdot \frac{\partial \ln \Gamma(w)}{\partial \ln w} \cdot d \ln w. \quad (\text{A56})$$

Stacking these equations and solving for wages, we obtain

$$d \ln w = \Theta^* \cdot \text{stack} \left(d \ln y + \sum_i \omega_{gi} \cdot z_{gi} + (\lambda - \eta) \cdot \sum_i \omega_{gi} \cdot d \ln p_i \right),$$

as claimed. \square

Appendix E Derivations for the Allen-Uzawa elasticities of substitution and properties of the propagation matrix

This section proves several properties of task shares, elasticities of substitution, and the propagation matrix mentioned in the text.

Symmetry of the task-share Jacobian Eq. (3) shows that the task-share Jacobian satisfies a symmetry property. To prove this, consider a proportional increase in w_g by $\Delta w_g = w_g \cdot \epsilon$ for some $\epsilon > 0$, a set $\mathcal{M}(\epsilon)$ of these tasks are assigned to g' and increase g' task share by $\Delta \Gamma_{g'} = \int_{\mathcal{M}(\epsilon)} \psi_{g'}(x)^{\lambda-1} \cdot dx$. Therefore,

$$\frac{\partial \Gamma_{g'}(w)}{\partial w_g} = \lim_{\epsilon \rightarrow 0} \frac{\int_{\mathcal{M}(\epsilon)} \psi_{g'}(x)^{\lambda-1} \cdot dx}{w_g \cdot \epsilon}.$$

Now, suppose that $w_{g'}$ decreases proportionally by $\Delta w_{g'} = -w_{g'} \cdot \epsilon$ for some $\epsilon > 0$. The same set $\mathcal{M}(\epsilon)$ of tasks switch to g' and decrease skill group g 's task share by $\Delta \Gamma_g = -\int_{\mathcal{M}(\epsilon)} \psi_g(x)^{\lambda-1} \cdot dx$. Now noting that for marginal tasks we have $\frac{w_g}{A_g \cdot \psi_g(x)} = \frac{w'_g}{A_{g'} \cdot \psi_{g'}(x)}$, we can conclude

$$\begin{aligned} \frac{\partial \Gamma_g(w)}{\partial w_{g'}} &= \lim_{\epsilon \rightarrow 0} \frac{\int_{\mathcal{M}(\epsilon)} \psi_{g'}(x)^{\lambda-1} \cdot \left(\frac{w_g}{w'_g} \right)^{\lambda-1} \cdot \left(\frac{A_{g'}}{A_g} \right)^{\lambda-1} \cdot dx}{w_{g'} \cdot \epsilon} \\ &= \left(\frac{w_g}{w'_g} \right)^\lambda \cdot \left(\frac{A_{g'}}{A_g} \right)^{\lambda-1} \cdot \frac{\partial \Gamma_{g'}(w)}{\partial w_g}. \end{aligned}$$

Properties of the propagation matrix We now prove the properties of the propagation matrix mentioned for the one-sector economy.

I. Dampening: Gershgorin's circle theorem in the proof of Lemma A1 already implied that the real part of all eigenvalues of \mathcal{H} are above 1. We now show that all eigenvalues of \mathcal{H} are real. To show this, first note that $\text{diag}(s^y) \mathcal{H} = \mathcal{H}_{\text{sym}}$ is a symmetric matrix with off-diagonal entry gg' given by $-\frac{1}{\lambda} \cdot s_g^y \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_j}$

and entry $g'g$ given by $-\frac{1}{\lambda} \cdot s_{g'}^y \cdot \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_g}$, which are equal due to the symmetry property of the Jacobian. Suppose ζ is an eigenvalue of \mathcal{H} with eigenvector v . Using upper bars to denote complex conjugates and superscript T to denote the transpose operation, we obtain

$$\begin{aligned}
\zeta \cdot \bar{v}^T \cdot \text{diag}(s^y) \cdot v &= \bar{v}^T \cdot (\text{diag}(s^y) \cdot \zeta \cdot v) \\
&= \bar{v}^T \cdot (\text{diag}(s^y) \cdot \mathcal{H} \cdot v) \\
&= \bar{v}^T \cdot (\mathcal{H}_{sym} \cdot v) \\
&= (\mathcal{H}_{sym} \cdot \bar{v})^T \cdot v \\
&= (\overline{\mathcal{H}_{sym}} \cdot \bar{v})^T \cdot v \\
&= (\text{diag}(s^y) \cdot \overline{\mathcal{H}} \cdot \bar{v})^T \cdot v \\
&= (\bar{\zeta} \cdot \text{diag}(s^y) \cdot \bar{v})^T \cdot v \\
&= \bar{\zeta} \cdot \bar{v}^T \cdot \text{diag}(s^y) \cdot v.
\end{aligned}$$

This string of identities implies that ζ equals its complex conjugate $\bar{\zeta}$ (since $v^T \cdot \text{diag}(s^y) \cdot v$ is a weighted vector norm, which must be positive) and must therefore be real. The justification for the steps involved is as follows. The first line uses the fact that ζ is a scalar. The second line uses the fact that ζ is an eigenvalue with eigenvector v . The third line uses the definition of \mathcal{H}_{sym} . The fourth line applies the transpose operator and uses the symmetry of \mathcal{H}_{sym} . The fifth line uses the fact that \mathcal{H}_{sym} is real. The sixth line uses once more the definition of \mathcal{H}_{sym} . The seventh line uses the fact that $\bar{\zeta}$ is also an eigenvalue of \mathcal{H}_{sym} with eigenvector \bar{v} . The last line applies the transpose operator once more. The idea behind the claim is intuitive: \mathcal{H} is a stretched version of a real symmetric matrix (which must therefore have all real eigenvalues and eigenvectors), and such stretching should not introduce complex eigenvalues.

The above derivations then show that all eigenvalues of \mathcal{H} are real and in $(1, \infty)$. This implies that all eigenvalues of $\Theta = \frac{1}{\lambda} \cdot \mathcal{H}^{-1}$ are also real and in $[0, 1/\lambda]$.

II. Monotonicity: We now turn to the monotonicity property, which says that $\theta_{gg} > \theta_{g'g}$ along a column. Suppose to obtain a contradiction that $\theta_{g'g} \geq \theta_{gg}$ and let $g' = \arg \max \theta_{g'g}$ be the index for the maximum along column g . We have that $\mathcal{H} \cdot \Theta = \frac{1}{\lambda}$. This requires entry $g'g$ in this product to be zero or

$$(1 - \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_{g'}}) \cdot \theta_{g'g} = \sum_{j \neq g', g} \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_j} \cdot \theta_{jg} + \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_g} \cdot \theta_{gg}.$$

By assumption, θ_{jg} and θ_{gg} are all less than or equal to $\theta_{g'g}$. This implies

$$(1 - \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_{g'}}) \cdot \theta_{g'g} \leq \sum_{j \neq g', g} \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_j} \cdot \theta_{g'g} + \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_g} \cdot \theta_{g'g},$$

dividing by $\theta_{g'g}$ and rearranging, we see that this yields

$$1 \leq \sum_j \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_j},$$

which is a contradiction since the sums $\sum_j \frac{1}{\lambda} \cdot \frac{\partial \ln \Gamma_{g'}(w)}{\partial \ln w_j}$ are 0 or negative (a common increase in wages causes all workers to loose tasks to capital).

III. Row sums: We now turn to the properties of the row sums of the propagation matrix, denoted by ρ_g . First, note that the elasticity of substitution between capital and group g can also be written in symmetrical form as

$$\sigma_{kg} = \sigma_{gk} = \lambda - \frac{1}{s_K^y} \cdot \sum_{g'} \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}},$$

since a percent increase in the user cost of capital generates the same substitution patterns as a commensurate percent reduction in all wages. This identity can be written in matrix form as

$$-\frac{1}{\lambda} \frac{\partial \ln \Gamma(w)}{\partial \ln w} \cdot \text{stack}(1) = \text{stack} \left(s_K^y \cdot \left(\frac{\sigma_{kg}}{\lambda} - 1 \right) \right),$$

or equivalently

$$\mathcal{H} \cdot \text{stack}(1) = \text{stack} \left(1 + s_K^y \cdot \left(\frac{\sigma_{kg}}{\lambda} - 1 \right) \right).$$

Multiplying by Θ on the left of both sides yields

$$\frac{1}{\lambda} \cdot \text{stack}(1) = \Theta \cdot \text{stack} \left(1 + s_K^y \cdot \left(\frac{\sigma_{kg}}{\lambda} - 1 \right) \right).$$

Comparing row g on both sides, we get

$$\rho_g + s_K^y \cdot \sum_{g'} \theta_{gg'} \cdot \left(\frac{\sigma_{kg'}}{\lambda} - 1 \right) = \frac{1}{\lambda},$$

or equivalently

$$\rho_g = \frac{1}{\lambda} \cdot \left[1 + s_K^y \cdot \left(\frac{\bar{\sigma}_{kg}}{\lambda} - 1 \right) \right]^{-1},$$

which gives the formula in the main text. Note that this formula implies that $\rho_g \in (0, 1/\lambda]$, as also claimed in the main text.

IV. Relationship to elasticities of substitution: We now derive the expression that relates the propagation matrix to the matrix of elasticities of substitution Σ .

First, we have

$$\sigma_{gg} = \frac{1}{s_g^y} \cdot \frac{d \ln \ell_g}{d \ln w_g} \Bigg|_{y \text{ constant}} = \lambda - \frac{\lambda}{s_g^y} + \frac{1}{s_g^y} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_g},$$

$$\sigma_{gg'} = \frac{1}{s_{g'}^y} \cdot \frac{d \ln \ell_g}{d \ln w_{g'}} \Bigg|_{y \text{ constant}} = \lambda + \frac{1}{s_{g'}^y} \cdot \frac{\partial \ln \Gamma_g(w)}{\partial \ln w_{g'}}.$$

We can then write

$$\Sigma = \lambda - \lambda \cdot \text{diag} \left(\frac{1}{s^y} \right) + \frac{\partial \ln \Gamma}{\partial \ln w} \cdot \text{diag} \left(\frac{1}{s^y} \right).$$

Rearranging this yields

$$\mathcal{H} \cdot \lambda \cdot \text{diag} \left(\frac{1}{s^y} \right) = \lambda - \Sigma.$$

Pre-multiplying by Θ on both sides yields

$$\text{diag} \left(\frac{1}{s^y} \right) = \Theta \cdot (\lambda - \Sigma),$$

and solving for Θ yields the relationship outlined in the text

$$\Theta = \text{diag} \left(\frac{1}{s^y} \right) \cdot (\lambda - \Sigma)^{-1}.$$

V Symmetry: The above identity also guarantees that $\text{diag}(s^y) \cdot \Theta = (\lambda - \Sigma)^{-1}$ is symmetric, which implies $\theta_{gg'}/s_{g'}^y = \theta_{g'g}/s_g^y$.

Appendix F Additional empirical results

F.1 Robustness checks

The tables in this part of the Appendix report a series of robustness checks on our reduced-form analysis.

- Table A1 reports the same specifications shown in Table 1 for wages in the main text, but proxies for new tasks as

$$\begin{aligned} d \ln \Gamma_g^{\text{new}} &= \sum_o \omega_{go}^{1980} \cdot \text{Share new job titles DOT 1977} \\ &+ \sum_o \omega_{go}^{1980} \cdot \text{Share new job titles DOT 1991} \\ &+ \sum_o \omega_{go}^{1980} \cdot \text{Share new job titles Census 2000}. \end{aligned}$$

This measure apportions new tasks across groups based on 1980 employment shares.

- Table A2 reports the same specifications shown in Table 2 for hours worked per person in the main text, but apportions new tasks across groups based on 1980 employment shares.
- Table A3 decomposes the effects of automation and new tasks into an extensive and intensive margin of employment.
- Table A4 reports estimates for wages and hours worked separately for workers with no college degree and those with a college degree.

F.2 Estimating the propagation matrix

Once we impose our parameterization of the Jacobian, we can rewrite the estimating equation in (44) as

$$\begin{aligned}
 \sigma \Delta \ln w_g + d \ln \Gamma_g^{\text{auto}} - d \ln \Gamma_g^{\text{new}} \\
 = \tilde{\beta} X_g + \gamma \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) \\
 + \gamma_{\text{job}} \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{job similarity}_{gg'} \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) \\
 + \gamma_{\text{edu-age}} \cdot \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{edu-age similarity}_{gg'} \cdot (\Delta \ln w_{g'} - \Delta \ln w_g) \\
 + \tilde{\nu},
 \end{aligned}$$

where $\tilde{\beta}$ and $\tilde{\nu}$ are linear transformations of β and ν respectively.

This equation can be estimated via GMM/2SLS after imposing $\sigma = \lambda + \varphi = 0.6$ (as discussed in the text). Our estimation imposes the restriction that $\gamma, \gamma_{\text{job}}, \gamma_{\text{edu-age}} \geq 0$.

The ripple terms on the right hand side are instrumented using

$$\begin{aligned}
 Z_g &= \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g) \\
 Z_{\text{job},g} &= \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{job similarity}_{gg'} \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g) \\
 Z_{\text{edu-age},g} &= \sum_{g'} \sum_n \omega_{gn} \cdot s_{g'}^n \cdot \text{edu-age similarity}_{gg'} \cdot (\Delta \ln \hat{w}_{g'} - \Delta \ln \hat{w}_g),
 \end{aligned}$$

respectively. Here $\Delta \ln \hat{w}_g$ is the predicted wage change based on groups experienced task displacement from automation, exposure to new tasks, and the exogenous covariates in the model. We get very similar results if we instead use $\Delta \ln \hat{w}_g = d \ln \Gamma_g^{\text{new}} - d \ln \Gamma_g^{\text{auto}}$ to form these instruments.

TABLE A1 Reduced-form evidence: changes in real hourly wages regressed on automation and new tasks, 1980–2016. Robustness check using alternative measure of new tasks.

	Dependent variables: Change in log hourly wages, 1980–2016						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Only displacement from automation							
Automation task displacement	−1.65 (0.10)	−1.41 (0.20)	−1.50 (0.11)	−1.45 (0.18)	−1.41 (0.19)	−1.71 (0.25)	−1.75 (0.32)
R^2 for model	0.64	0.66	0.69	0.82	0.83	0.76	0.76
R^2 for automation	0.64	0.55	0.59	0.56	0.55	0.67	0.68
R^2 remaining covs		0.11	0.10	0.26	0.28	0.09	0.08
Observations	500	500	500	500	500	492	492
Panel B. Only reinstatement from new tasks							
New tasks reinstatement	2.82 (0.21)	3.57 (0.45)	3.07 (0.25)	2.93 (0.46)	2.52 (0.52)	3.39 (0.62)	3.47 (0.82)
R^2 for model	0.63	0.64	0.65	0.79	0.79	0.65	0.58
R^2 for new tasks	0.63	0.80	0.69	0.66	0.56	0.76	0.78
R^2 remaining covs		−0.16	−0.04	0.14	0.22	−0.11	−0.19
Observations	500	500	500	500	500	492	492
Panel C. Both explanatory variables							
Automation task displacement	−0.94 (0.26)	−0.90 (0.26)	−1.05 (0.26)	−1.17 (0.27)	−1.25 (0.26)	−1.42 (0.31)	−1.53 (0.31)
New tasks reinstatement	1.46 (0.47)	2.06 (0.61)	1.13 (0.54)	1.09 (0.75)	0.72 (0.71)	0.98 (0.79)	0.80 (0.76)
R^2 for model	0.69	0.70	0.70	0.83	0.83	0.78	0.77
R^2 for automation	0.37	0.35	0.41	0.46	0.49	0.55	0.60
R^2 for new tasks	0.33	0.46	0.25	0.24	0.16	0.22	0.18
R^2 remaining covs		−0.11	0.04	0.13	0.18	0.00	−0.01
Observations	500	500	500	500	500	492	492
Panel D. Net task change due to new tasks minus automation							
Net task change (new tasks−automation)	1.12 (0.06)	1.18 (0.05)	1.07 (0.07)	1.15 (0.13)	1.12 (0.15)	1.31 (0.18)	1.35 (0.24)
R^2 for model	0.69	0.69	0.70	0.83	0.83	0.78	0.77
R^2 for automation	0.69	0.72	0.66	0.71	0.69	0.80	0.83
R^2 remaining covs		−0.03	0.04	0.12	0.14	−0.02	−0.06
Observations	500	500	500	500	500	492	492
<i>Other covariates:</i>							
Sectoral value added		✓		✓		✓	
Sectoral TFP			✓		✓		✓
Sectoral markups		✓		✓		✓	
Gender and education dummies			✓	✓	✓	✓	
Labor supply shifts					✓	✓	

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hourly wages across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The specifications are the same as in Table 1. The difference is that we now use a measure of new tasks that holds occupational shares fixed in 1980. The dependent variable is the change in log hourly wages for each group between 1980 and 2016. Panel A reports results using only our task displacement measure. Panel B only uses our task reinstatement measure. Panel C includes both task displacement and task reinstatement on the right-hand side. Panel D combines task displacement and reinstatement into a single net task change measure. The bottom rows list additional covariates included in each specification. As in Acemoglu and Restrepo (2022), we instrument changes in labor supply in columns 6 and 7 using trends in total hours worked by group from 1970 to 1980. All regressions are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

TABLE A2 Reduced-form evidence: changes in hours worked per person regressed on automation and new tasks, 1980–2016. Robustness check using alternative measure of new tasks.

	Dependent variables: Change in log hours worked per person, 1980–2016						
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
Panel A. Only displacement from automation							
Automation task displacement	−2.25 (0.30)	−1.58 (0.40)	−1.96 (0.27)	−1.83 (0.40)	−1.93 (0.41)	−2.21 (0.61)	−2.59 (0.78)
R^2 for model	0.44	0.48	0.50	0.68	0.67	0.61	0.56
R^2 for automation	0.44	0.31	0.38	0.36	0.38	0.43	0.51
R^2 remaining covs		0.17	0.11	0.32	0.29	0.18	0.05
Observations	500	500	500	500	500	492	492
Panel B. Only reinstatement from new tasks							
New tasks reinstatement	4.47 (0.53)	6.15 (1.21)	4.84 (0.50)	4.29 (0.99)	4.04 (1.01)	4.84 (1.32)	5.60 (1.58)
R^2 for model	0.59	0.61	0.60	0.68	0.65	0.57	0.43
R^2 for new tasks	0.59	0.81	0.64	0.56	0.53	0.64	0.74
R^2 remaining covs		−0.20	−0.04	0.11	0.12	−0.07	−0.30
Observations	500	500	500	500	500	492	492
Panel C. Both explanatory variables							
Automation task displacement	−0.22 (0.52)	−0.10 (0.52)	0.01 (0.48)	−1.25 (0.57)	−1.50 (0.60)	−1.56 (0.68)	−2.06 (0.89)
New tasks reinstatement	4.16 (1.04)	5.98 (1.64)	4.87 (0.93)	2.34 (1.47)	1.86 (1.52)	2.19 (1.47)	2.02 (1.56)
R^2 for model	0.59	0.61	0.60	0.69	0.67	0.64	0.58
R^2 for automation	0.04	0.02	−0.00	0.25	0.30	0.31	0.40
R^2 for new tasks	0.55	0.79	0.64	0.31	0.24	0.29	0.27
R^2 remaining covs		−0.20	−0.04	0.14	0.13	0.05	−0.09
Observations	500	500	500	500	500	492	492
Panel D. Net task change due to new tasks minus automation							
Net task change (new tasks-automation)	1.62 (0.20)	1.51 (0.34)	1.49 (0.18)	1.52 (0.29)	1.59 (0.29)	1.73 (0.44)	2.05 (0.57)
R^2 for model	0.53	0.53	0.55	0.69	0.67	0.64	0.58
R^2 for task changes	0.53	0.50	0.49	0.50	0.52	0.57	0.67
R^2 remaining covs		0.04	0.06	0.19	0.15	0.07	−0.09
Observations	500	500	500	500	500	492	492
<i>Other covariates:</i>							
Sectoral value added		✓		✓		✓	
Sectoral TFP			✓		✓		✓
Sectoral markups		✓		✓		✓	✓
Gender and education dummies				✓	✓	✓	✓
Labor supply shifts						✓	✓

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hours worked per person across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The specifications are the same as in Table 2. The difference is that we use a measure of new tasks that holds occupational shares fixed in 1980. The dependent variable is the change in log hours per person for each group between 1980 and 2016. Panel A reports results using only our task displacement measure. Panel B only uses our task reinstatement measure. Panel C includes both task displacement and task reinstatement on the right-hand side. Panel D combines task displacement and reinstatement into a single net task change measure. The bottom rows list additional covariates included in each specification. As in Acemoglu and Restrepo (2022), we instrument changes in labor supply in columns 6 and 7 using trends in total hours worked by group from 1970 to 1980. All regressions are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

TABLE A3 Reduced-form evidence: changes in hours intensive and extensive margin regressed on automation and new tasks, 1980–2016.

	Dependent variables:			
	Change in (log) employment to population ratios, 1980–2016		Change in (log) hours per working adult, 1980–2016	
	(1)	(2)	(3)	(4)
Automation task displacement	−0.77 (0.26)	−0.76 (0.26)	−0.99 (0.32)	−1.16 (0.31)
R^2 for model	0.73	0.72	0.42	0.42
R^2 for automation	0.19	0.18	0.34	0.40
R^2 remaining covs	0.54	0.54	0.08	0.02
Observations	500	500	500	500
Panel A. Only displacement from automation				
New tasks reinstatement	0.80 (0.41)	0.81 (0.48)	0.49 (0.48)	0.83 (0.56)
R^2 for model	0.71	0.71	0.35	0.35
R^2 for new tasks	0.16	0.16	0.11	0.18
R^2 remaining covs	0.55	0.54	0.25	0.17
Observations	500	500	500	500
Panel B. Only reinstatement from new tasks				
Automation task displacement	−0.70 (0.25)	−0.71 (0.25)	−0.99 (0.33)	−1.12 (0.31)
New tasks reinstatement	0.47 (0.35)	0.60 (0.44)	0.03 (0.42)	0.51 (0.50)
R^2 for model	0.73	0.72	0.42	0.42
R^2 for automation	0.17	0.17	0.34	0.39
R^2 for new tasks	0.10	0.12	0.01	0.11
R^2 remaining covs	0.47	0.43	0.07	−0.07
Observations	500	500	500	500
Panel C. Both explanatory variables				
Automation task displacement	−0.70 (0.25)	−0.71 (0.25)	−0.99 (0.33)	−1.12 (0.31)
New tasks reinstatement	0.47 (0.35)	0.60 (0.44)	0.03 (0.42)	0.51 (0.50)
R^2 for model	0.73	0.72	0.42	0.42
R^2 for automation	0.17	0.17	0.34	0.39
R^2 for new tasks	0.10	0.12	0.01	0.11
R^2 remaining covs	0.47	0.43	0.07	−0.07
Observations	500	500	500	500
Panel D. Net task change due to new tasks minus automation				
Net task change (new tasks-automation)	0.63 (0.20)	0.68 (0.21)	0.71 (0.24)	0.97 (0.25)
R^2 for model	0.73	0.72	0.41	0.42
R^2 for task changes	0.28	0.31	0.40	0.55
R^2 remaining covs	0.45	0.42	0.01	−0.13
Observations	500	500	500	500
<i>Other covariates:</i>				
Sectoral value added	✓		✓	
Sectoral TFP		✓		✓
Sectoral markups		✓		✓
Gender and education dummies	✓	✓	✓	✓

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hours worked per person across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The dependent variable is the change in (log) hours per worker (columns 1 and 2) and the change in (log) employment to population for each group between 1980 and 2016. Panel A reports results using only our task displacement measure. Panel B only uses our task reinstatement measure. Panel C includes both task displacement and task reinstatement on the right-hand side. Panel D combines task displacement and reinstatement into a single net task change measure. The bottom rows list additional covariates included in each specification. All regressions are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

TABLE A4 Reduced-form evidence: changes in real hourly wages and hours worked regressed on automation and new tasks, 1980–2016. Robustness check reporting estimates for groups with and without a college degree.

	Dependent variables:					
	Change (log) hourly wages, 1980–2016			Change (log) hours worked, 1980–2016		
	(1)	(2)	(3)	(4)	(5)	(6)
Panel A. Workers with no college degree						
Automation task displacement	−0.76 (0.31)	−1.16 (0.20)	−1.20 (0.22)	−1.12 (0.57)	−1.59 (0.50)	−1.74 (0.53)
New tasks reinstatement	1.04 (0.42)	2.16 (0.76)	1.95 (0.79)	3.47 (0.87)	2.97 (1.47)	2.07 (1.91)
R^2 for model	0.42	0.74	0.72	0.52	0.64	0.61
R^2 for automation	0.30	0.45	0.47	0.22	0.31	0.34
R^2 for new tasks	0.25	0.52	0.47	0.49	0.42	0.29
R^2 remaining covs		−0.24	−0.21		−0.08	−0.03
Observations	300	300	300	300	300	300
Panel B. Workers with a college degree						
Automation task displacement	−2.34 (0.58)	−1.84 (0.62)	−1.56 (0.49)	−0.87 (0.80)	−2.14 (0.78)	−1.16 (0.70)
New tasks reinstatement	0.86 (0.28)	0.83 (0.21)	0.93 (0.25)	−0.11 (0.37)	0.07 (0.34)	0.20 (0.42)
R^2 for model	0.21	0.60	0.59	0.03	0.64	0.60
R^2 for automation	0.91	0.72	0.61	0.17	0.42	0.23
R^2 for new tasks	0.21	0.20	0.22	−0.01	0.01	0.03
R^2 remaining covs		−0.32	−0.25		0.21	0.34
Observations	200	200	200	200	200	200
<i>Other covariates:</i>						
Sectoral value added		✓			✓	
Sectoral TFP			✓			✓
Sectoral markups			✓			✓
Gender and education dummies		✓	✓		✓	✓

Notes: This table presents estimates of the relationship between automation, new tasks, and the change in hourly wages and hours worked per person across 500 demographic groups, defined by gender, education, age, race, and native/immigrant status. The dependent variable is the change in (log) hourly wages (columns 1–3) and the change in (log) hours worked (columns 4–6) from 1980 and 2016. Panel A provides estimates for groups of workers with no college degree. Panel B provides estimates for groups of workers with a college degree. The bottom rows list additional covariates included in each specification. All regressions are weighted by total hours worked by each group in 1980. Standard errors robust to heteroskedasticity are reported in parentheses.

References

Acemoglu, Daron, 2002. Technical change, inequality, and the labor market. *Journal of Economic Literature* 40, 7–72.

Acemoglu, Daron, 2024. The Simple Macroeconomics of AI. Tech. Rep.

Acemoglu, Daron, Autor, David, 2011. Skills, tasks and technologies: implications for employment and earnings. In: *Handbook of Labor Economics*, vol. 4. Elsevier, pp. 1043–1171. Chap. 12.

Acemoglu, Daron, Autor, David, Hazell, Jonathon, Restrepo, Pascual, 2022. Artificial intelligence and jobs: evidence from online vacancies. *Journal of Labor Economics* 40, S293–S340.

Acemoglu, Daron, Gancia, Gino, Zilibotti, Fabrizio, 2015. Offshoring and directed technical change. *American Economic Journal: Macroeconomics* 7, 84–122.

Acemoglu, Daron, Jensen, Martin Kaae, Restrepo, Pascual, 2024. When Does Technology Increase Long-Run Wages? Tech. Rep. Massachusetts Institute of Technology.

Acemoglu, Daron, Koster, Hans R.A., Ozgen, Ceren, 2023. Robots and Workers: Evidence from the Netherlands. Working Paper 31009. National Bureau of Economic Research.

Acemoglu, Daron, Lelarge, Claire, Restrepo, Pascual, 2020. Competing with robots: firm-level evidence from France. *AEA Papers and Proceedings* 110, 383–388.

Acemoglu, Daron, Loebbing, Jonas, 2024. Automation and Polarization. Tech. Rep.

Acemoglu, Daron, Restrepo, Pascual, 2018a. Artificial intelligence, automation, and work. In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, pp. 197–236.

Acemoglu, Daron, Restrepo, Pascual, 2018b. The race between man and machine: implications of technology for growth, factor shares, and employment. *American Economic Review* 108, 1488–1542.

Acemoglu, Daron, Restrepo, Pascual, 2019. Automation and new tasks: how technology displaces and reinstates labor. *The Journal of Economic Perspectives* 33, 3–30.

Acemoglu, Daron, Restrepo, Pascual, 2020a. Robots and jobs: evidence from US labor markets. *Journal of Political Economy* 128, 2188–2244.

Acemoglu, Daron, Restrepo, Pascual, 2020b. Unpacking skill bias: automation and new tasks. *AEA Papers and Proceedings* 110, 356–361.

Acemoglu, Daron, Restrepo, Pascual, 2022. Tasks, automation, and the rise in US wage inequality. *Econometrica* 90, 1973–2016.

Acemoglu, Daron, Restrepo, Pascual, 2024. Automation and Rent Dissipation: Implications for Inequality, Productivity, and Welfare. Tech. Rep.

Acemoglu, Daron, Zilibotti, Fabrizio, 2001. Productivity differences. *The Quarterly Journal of Economics* 116, 563–606.

Aghion, Philippe, Jones, Benjamin F., Jones, Charles I., 2018. Artificial intelligence and economic growth. In: *The Economics of Artificial Intelligence: An Agenda*. University of Chicago Press, pp. 237–282.

Autor, David, Chin, Caroline, Salomons, Anna, Seegmiller, Bryan, 2024. New frontiers: the origins and content of New Work, 1940–2018. *The Quarterly Journal of Economics* 139 (3), 1399–1465.

Autor, David H., 2019. Work of the past, work of the future. *AEA Papers and Proceedings* 109, 1–32.

Autor, David H., Handel, Michael J., 2013. Putting tasks to the test: human capital, job tasks, and wages. *Journal of Labor Economics* 31, S59–S96.

Autor, David H., Katz, L.F., Krueger, A.B., 1998. Computing inequality: have computers changed the labor market? *The Quarterly Journal of Economics* 113, 1169–1213.

Autor, David H., Levy, Frank, Murnane, Richard J., 2003. The skill content of recent technological change: an empirical exploration. *The Quarterly Journal of Economics* 118, 1279–1333.

Babina, Tania, Fedyk, Anastassia, He, Alex, Hodson, James, 2024. Artificial intelligence, firm growth, and product innovation. *Journal of Financial Economics* 151, 103745.

Baumol, William J., de Ferranti, David, Malach, Monte, Pablos-Méndez, Ariel, Tabish, Hilary, Wu, Lilian Gomory, 2012. *The Cost Disease: Why Computers Get Cheaper and Health Care Doesn't*. Yale University Press.

Berman, Eli, Bound, John, Griliches, Zvi, 1994. Changes in the demand for skilled labor within U. S. manufacturing: evidence from the annual survey of manufactures. *The Quarterly Journal of Economics* 109, 367–397.

Bonfiglioli, Alessandra, Crinò, Rosario, Fadinger, Harald, Gancia, Gino, 2020. Robot Imports and Firm-Level Outcomes. Tech. Rep., CEPR Discussion Papers No. 14593.

Bound, John, Johnson, George, 1992. Changes in the structure of wages in the 1980's: an evaluation of alternative explanations. *American Economic Review* 82, 371–392.

Boustan, Leah Platt, Choi, Jiwon, Clingingsmith, David, 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., Kaboski, Joseph P., Rogerson, Richard, Vizcaino, Juan I., 2021. Skill-biased structural change. *The Review of Economic Studies* 89, 592–625.

Cahuc, Pierre, 2024. Labour market inequality in France and in the UK. *Oxford Open Economics* 3, i1933–i1939.

Card, David, Lemieux, Thomas, 2001. Can falling supply explain the rising return to college for younger men? A cohort-based analysis. *The Quarterly Journal of Economics* 116, 705–746.

Caselli, Francesco, Manning, Alan, 2019. Robot arithmetic: new technology and wages. *American Economic Review: Insights* 1, 1–12.

Chetty, Raj, Guren, Adam, Manoli, Day, Weber, Andrea, 2011. Are micro and macro labor supply elasticities consistent? A review of evidence on the intensive and extensive margins. *American Economic Review* 101, 471–475.

De Loecker, Jan, Eeckhout, Jan, Unger, Gabriel, 2020. The rise of market power and the macroeconomic implications. *The Quarterly Journal of Economics* 135, 561–644.

Dechezleprêtre, Antoine, Hémous, David, Olsen, Morten, Zanella, Carlo, 2023. Induced Automation Innovation: Evidence from Firm-level Patent Data. Working Paper. University of Zurich, Department of Economics.

Dornbusch, Rüdiger, Fischer, Stanley, Samuelson, Paul, 1977. Comparative advantage, trade, and payments in a Ricardian model with a continuum of goods. *American Economic Review* 67, 823–839.

Dustmann, Christian, Ludsteck, Johannes, Schönberg, Uta, 2009. Revisiting the German wage structure. *The Quarterly Journal of Economics* 124, 843–881.

Dvorkin, Maximiliano A., Monge-Naranjo, Alexander, 2019. Occupation mobility, human capital and the aggregate consequences of task-biased innovations. FRB St. Louis Working Paper.

Eaton, Jonathan, Kortum, Samuel, 2002. Technology, geography, and trade. *Econometrica* 70, 1741–1779.

Feng, Andy, Graetz, Georg, 2020. Training requirements, automation, and job polarisation. *The Economic Journal* 130, 2249–2271.

Freund, Lukas, 2024. Superstar Teams. Tech. Rep. Cambridge University.

Fuchs-Schündeln, Nicola, Krueger, Dirk, Sommer, Mathias, 2010. Inequality trends for Germany in the last two decades: a tale of two countries. *Review of Economic Dynamics* 13, 103–132.

Gale, David, Nikaido, Hukukane, 1965. The Jacobian matrix and global univalence of mappings. *Mathematische Annalen* 159, 81–93.

Goldin, Claudia Dale, Katz, Lawrence F., 2008. The Race Between Education and Technology. Harvard University Press, Cambridge.

Goos, Maarten, Manning, Alan, 2007. Lousy and lovely jobs: the rising polarization of work in Britain. *Review of Economics and Statistics* 89, 118–133.

Goos, Maarten, Manning, Alan, Salomons, Anna, 2014. Explaining job polarization: routine-biased technological change and offshoring. *American Economic Review* 104, 2509–2526.

Gornick, Janet C., 2024. Income inequality and income poverty in a cross-national perspective. *Oxford Open Economics* 3, i147–i155.

Graetz, Georg, Michaels, Guy, 2018. Robots at work. *Review of Economics and Statistics* 100, 753–768.

Griliches, Zvi, 1969. Capital-skill complementarity. *Review of Economics and Statistics* 51, 465–468.

Grossman, Gene M., Rossi-Hansberg, Esteban, 2008. Trading tasks: a simple theory of offshoring. *American Economic Review* 98, 1978–1997.

Hubmer, Joachim, Restrepo, Pascual, 2021. Not a Typical Firm: The Joint Dynamics of Firms, Labor Shares, and Capital-Labor Substitution. Working Paper 28579. National Bureau of Economic Research.

Hummel, Anders, 2020. Robot Adoption and Labor Market Dynamics. Working paper. University of Chicago.

Jones, Benjamin F., Liu, Xiaojie, 2022. A Framework for Economic Growth with Capital-Embodied Technical Change. Working Paper 30459. National Bureau of Economic Research.

Karabarbounis, Loukas, Neiman, Brent, 2013. The global decline of the labor share. *The Quarterly Journal of Economics* 129, 61–103.

Katz, Lawrence F., Autor, David H., 1999. Changes in the wage structure and earnings inequality. In: *Handbook of Labor Economics*, vol. 3. Elsevier, pp. 1463–1555.

Katz, Lawrence F., Murphy, Kevin M., 1992. Changes in relative wages, 1963–1987: supply and demand factors. *The Quarterly Journal of Economics* 107, 35–78.

Kikuchi, Shinnosuke, 2024. Automation, Fragmentation, and Changes in Comparative Advantage. Tech. Rep.

Kim, Ryan, Vogel, Jonathan, 2021. Trade shocks and labor market adjustment. *American Economic Review: Insights* 3, 115–130.

Klenow, Peter J., Rodríguez-Clare, Andrés, 1997. The neoclassical revival in growth economics: has it gone too far? *NBER Macroeconomics Annual* 12, 73–103.

Kogan, Leonid, Papanikolaou, Dimitris, Schmidt, Lawrence D.W., Seegmiller, Bryan, 2021. Technology, Vintage-Specific Human Capital, and Labor Displacement: Evidence from Linking Patents with Occupations. Working Paper 29552. National Bureau of Economic Research.

Krueger, Dirk, Perri, Fabrizio, Pistaferri, Luigi, Violante, Giovanni L., 2010. Cross-sectional facts for macroeconomists. *Review of Economic Dynamics* 13, 1–14.

Krusell, Per, Ohanian, Lee E., Ríos-Rull, José-Víctor, Violante, Giovanni L., 2000. Capital-skill complementarity and inequality: a macroeconomic analysis. *Econometrica* 68, 1029–1053.

Lin, Jeffrey, 2011. Technological adaptation, cities, and new work. *Review of Economics and Statistics* 93, 554–574.

Lind, Nelson, Ramondo, Natalia, 2023. Trade with correlation. *American Economic Review* 113, 317–353.

Machin, Stephen, Van Reenen, John, 2010. Inequality: Still Higher, but Labour's Policies Kept It down. Tech. Rep. 15. Center for Economic Performance, The London School of Economics and Political Science.

Moll, Benjamin, Rachel, Lukasz, Restrepo, Pascual, 2022. Uneven growth: automation's impact on income and wealth inequality. *Econometrica* 90, 2645–2683.

Nakamura, Hideki, Zeira, Joseph, 2024. Automation and unemployment: help is on the way. *Journal of Economic Growth* 29, 215–250.

Ngai, L. Rachel, Pissarides, Christopher A., 2007. Structural change in a multisector model of growth. *American Economic Review* 97, 429–443.

Oberfield, Ezra, Raval, Devesh, 2021. Micro data and macro technology. *Econometrica* 89, 703–732.

Peretto, Pietro F., Seater, John J., 2013. Factor-eliminating technical change. *Journal of Monetary Economics* 60, 459–473.

Restrepo, Pascual, 2024. Automation: theory, evidence, and outlook. *Annual Review of Economics* 16, 1–25.

Rodríguez-Clare, Andrés, 2010. Offshoring in a Ricardian world. *American Economic Journal: Macroeconomics* 2, 227–258.

Sattiger, Michael, 1975. Comparative advantage and the distributions of earnings and abilities. *Econometrica* 43, 455–468.

Simon, Herbert Alexander, 1965. *The Shape of Automation for Men and Management*. vol. 13. Harper & Row, New York.

Teulings, Coen N., 1995. The wage distribution in a model of the assignment of skills to jobs. *Journal of Political Economy* 103, 280–315.

Teulings, Coen N., 2005. Comparative advantage, relative wages, and the accumulation of human capital. *Journal of Political Economy* 113, 425–461.

Van Reenen, John, 2011. Wage inequality, technology and trade: 21st century evidence. *Labour Economics* 18, 730–741.

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.