PSL Quarterly Review

vol. 79 no. 318 (September 2026)

Article

How should macroeconomists think about AI and the labor market?

EDOARDO GAFFEO*

Abstract:
How does generative artificial intelligence (AI) affect labor markets? Existing models of skill-biased and routine-biased technical change do not fully explain recent worker-level evidence: early employment effects are concentrated among highly educated, entry-level workers in prediction-intensive occupations, appear chiefly in hiring rather than wages, and remain limited in the aggregate. This paper develops a conceptual framework that integrates task-based automation with a decision-theoretic account of AI as cognitive augmentation. AI reshapes the allocation of tasks by changing how firms combine prediction, judgment, and expertise. The framework connects these organizational adjustments to aggregate labor-market outcomes and sets out a research agenda on AI's effects on employment, wage inequality, and economic growth.

University of Trento, Italy, email: edoardo.gaffeo@unitn.it

How to cite this article:
Gaffeo E. (2026), "How should macroeconomists think about AI and the labor market?", *PSL Quarterly Review*, 79 (318), pp. 221-251.

DOI: https://doi.org/10.13133/2037-3643/19268

JEL codes:
E24, J24, O33

Keywords:
artificial intelligence, task-based models, cognitive models, aggregate labor market

Journal homepage:
https://rosa.uniroma1.it/rosa04/psl_quarterly_review

Over the last two decades, artificial intelligence (AI) – defined as machine-based systems that utilize data-driven learning to generate predictions and decisions – has moved from a specialized subfield of computer science to a central stage in macroeconomic inquiry. Unlike previous waves of automation associated e.g. to internal combustion engine and electricity, which primarily displaced routine manual labor while leaving the employment-population ratio almost unaffected until the 2000s (Moffitt, 2012; Autor and Solomons, 2018; Abraham and Kearney, 2020; see also Felipe et al., 2020, for a critical exposition), the rapid diffusion of digitalization and recent breakthroughs in machine learning and computational data processing increasingly target cognitive domains once considered the exclusive province of white-collar professionals. As a result, deep concerns surrounding employment, career trajectories, and the distributional consequences of technological change have intensified across a broad range of jobs. These themes

* Many thanks to Lukas B. Freund, Lucio Gobbi, Roberto Tamborini, seminar participants to the "Renaissance in Economics" 2026 Conference, two anonymous referees and the Managing Editor for helpful comments on previous versions of the paper.
The author declares no personal or financial relationships that could be perceived as potential conflicts of interest with respect to the research, authorship, dissemination or publication of this paper.

This work is licensed under a Creative Commons Attribution – Non-Commercial – No Derivatives 4.0 International License. To view a copy of this license visit http://creativecommons.org/licenses/by-nc-nd/4.0/

have featured prominently in public discourse (see, for example, Cain Miller and Cox, 2023; *The Financial Times*, 2025) and have also become a focal point of a growing academic literature (Brynjolfsson and McAfee, 2014; Frey and Osborne, 2017; Korinek and Juelfs, 2022; see also the systematic review by Nigar et al., 2025).

While the claim that "this time is different" is usually met with disenchanted skepticism by economists, the case for generative AI's uniqueness is conceptually grounded. Earlier digital automation and robotics relied on explicit programming to mechanize well-defined tasks. In contrast, modern AI operates through data-driven inference, fundamentally altering the nature of human-machine interaction (Agrawal et al., 2018). Yet, despite a rich theoretical literature, empirical evidence on aggregate labor-market outcomes has remained fragmented. Most extant studies rely on static, task-based exposure measures – often derived from occupational surveys – that lack the longitudinal granularity required to track worker-level transitions. Furthermore, by focusing on partial-equilibrium effects within specific sectors, these studies often overlook the general-equilibrium adjustments and firm-level heterogeneities that dictate aggregate demand.

The fresh availability of granular, worker-level data spanning a sufficiently long period marks a substantive turning point in the empirical literature on AI and labor markets. Most notably, Brynjolfsson et al. (2025a; hereafter BCC) provide the first – and, admittedly, much debated – large-scale analysis based on administrative, high-frequency records that links measures of AI exposure to realized employment and earnings outcomes for several millions U.S. workers. Their evidence uncovers a set of striking and internally consistent patterns that materially help macroeconomists to refine the way they conceptualize the labor-market effects of AI. Three findings are particularly salient. First, exposure to AI is associated with meaningful reductions in employment, but these effects are sharply concentrated among younger and entry-level workers, including those with relatively high levels of formal human capital. However, the overall impact on aggregate employment is so far almost negligible.[1] Second, the adverse employment responses are largest in occupations characterized by intensive prediction and information-processing tasks, where accumulated job-specific experience appears to be complementary to AI adoption. Third, adjustment in response to AI exposure occurs predominantly along the extensive margin, with negligible effects on wages. Taken together, these results point to AI-induced changes operating primarily through a complex dance of task reallocations within occupations and firms, rather than through uniform shifts in relative factor prices, and act as a means to discipline theoretical explanations of the observed relationship between AI and employment. Although the extent to which these findings withstand competing explanations remains an open question – one to which we return below – they nonetheless reveal empirical regularities that are difficult to accommodate within canonical models of skill-biased technological change, as well as within task-based frameworks in which AI uniformly substitutes for labor across occupations. The emerging evidence instead points to the need for a framework capable of linking the microeconomics of AI adoption to aggregate labor market outcomes without abstracting from the heterogeneous ways in which AI reshapes production within firms.

The analysis that follows develops such a framework by integrating two influential, yet largely separate, strands of the literature. Building on Acemoglu and Restrepo (2018, 2019a, 2019b, 2022), production is viewed as a continuum of tasks over which technological change reallocates activities between labor and capital through an endogenously evolving automation frontier, generating displacement and reinstatement effects. From Agrawal et al. (2018, 2019a, 2019b, 2025, 2026) and Autor and Thompson (2025), we draw the complementary insight that AI

[1] Humlum and Vestergaard (2025) obtain a similar result with administrative data for Denmark.

reduces the cost of prediction and inference, thereby augmenting human cognition and reshaping the role of judgment within production. Bringing these perspectives together yields a framework in which AI operates simultaneously along three margins. It expands the range of tasks that can be performed autonomously by capital; it raises the effective productivity of workers in activities where judgment, interpretation, and contextual reasoning remain essential; and, by relaxing cognitive constraints, it enlarges the set of economically viable tasks in which labor can retain or acquire comparative advantage. Comparative advantage therefore evolves through the reallocation of existing tasks, changes in productivity within surviving human tasks, and the creation of new activities at the frontier of production. This interaction provides a unified interpretation of the emerging evidence on AI adoption while identifying the channels through which firm-level organizational change propagates into productivity growth, wage inequality, labor's share, and long-run economic growth.

The model is intended to organize, rather than settle, the debate by providing a framework within which the aggregate consequences of AI adoption can be studied systematically. It identifies a parsimonious set of economic mechanisms – and the structural parameters governing them – that link firm-level adoption to productivity growth, wage inequality, labor's share, and long-run economic growth. In doing so, it provides a common analytical language for interpreting the emerging microeconomic evidence while identifying the empirical objects most relevant for future macroeconomic research on AI.

This paper can also be read as a selective – though necessarily incomplete[2] – guide to a rapidly expanding literature that studies AI as a general-purpose technology[3] (Crafts, 2021; OECD, 2024) and, in doing so, extends the analysis of technological change beyond the traditional paradigms of skill-biased and routine-biased technological change. In the skill-biased technological change framework, technological progress disproportionately increases the productivity of more educated workers, raising the skill premium and widening wage inequality (Katz and Murphy, 1992; Goldin and Katz, 2008). The routine-biased technological change literature instead emphasizes the substitution of technology for routine tasks and its complementarity with abstract and non-routine activities, thereby accounting for the polarization of employment across occupations (Autor et al., 2003; Acemoglu and Autor, 2011). These frameworks successfully captured many of the defining features of earlier waves of computerization. The growing literature on AI, however, has increasingly emphasized dimensions that they treat only imperfectly, including the continuous reallocation of comparative advantage across tasks and the role of AI as a direct input into cognitive production. Against this background, the framework developed below uses labor market adjustment as an organizing lens through which to bring together insights from the emerging micro- and macroeconomics of AI. Its purpose is not to survey this literature exhaustively, but to identify the ideas that appear most relevant for understanding how AI adoption propagates through the economy and ultimately shapes productivity, employment, and the distribution of income.

The remainder of the paper proceeds as follows. Section 1 contextualizes our contribution within the empirical literature sparked by BCC, documenting the nuanced heterogeneity of generative AI's impact across tasks. Section 2 interprets these stylized facts through a conceptual

[2] For instance, I will not engage with questions concerning predicted TFP growth *via* Hulten theorem (Acemoglu, 2025), international trade (Jakubik et al., 2025) and global value chains (Filippucci et al., 2026), potential transformative effects of artificial general intelligence (Trammell and Korinek, 2025) like singularities in the growth process (Jones and Tonetti, 2026) or the emergence of existential risks (Jones, 2024).

[3] Following Bresnahan and Trajtenberg (1995), general-purpose technologies are characterized by pervasiveness, iterative improvement and the capacity for downstream innovation.

framework that explicitly delineates an "automation frontier" from an "augmentation margin". We formalize this framework in section 3, characterizing an integrated model of seniority-based technological change. Sections 4 and 5 deploy the model to examine its macroeconomic implications for wage inequality and balanced growth, respectively. Section 6 outlines a consistent research agenda for future empirical work. Section 7 concludes.

1. Granular evidence on AI exposure and labor market adjustments

A central challenge in the empirical analysis of technological change is establishing whether measures of ex ante exposure to new technologies translate into predictable adjustments in labor market outcomes. In the case of AI, this issue is particularly salient. A growing literature, often grounded in task-based frameworks, has characterized occupations according to their potential exposure to AI on the basis of the structure, content, and codifiability of their tasks (Acemoglu and Autor, 2011; Eloundou et al., 2024). By contrast, the vast majority of the available empirical evidence on the realized labor market effects of generative AI remains limited in scope and, to date, largely unsatisfying in deriving consolidated stylized facts (Acemoglu et al., 2022; Morandini et al., 2023; Pizzinelli et al., 2023; Poláková et al., 2023; Cazzaniga et al., 2024; Zarifhonarvar, 2024; Brynjolfsson et al., 2025b; Chanoi and Banger-Drowns, 2025; Goller et al., 2025). This reflects both the recency of technology's diffusion and the absence, in standard data sources, of the frequency and granularity required to detect early labor market adjustments.

BCC combined monthly data from administrative sources and payroll processing, covering between 3.5 and 5 million full-time workers per month and tens of thousands of U.S. firms and spanning from January 2021 to July 2025, to provide new evidence on this question. These data allow employment to be observed at the level of worker age, occupation, and firm, and substantially improve on traditional survey data in terms of sample size and temporal resolution. This feature is critical for identifying heterogeneous employment responses across narrowly defined demographic and occupational groups. Rather than studying long-run equilibrium adjustments, BCC focuses on short-run employment responses following the widespread availability of large language models (LLMs) in late 2022.[4]

To operationalize the concept of AI exposure of each occupation, BCC employ two complementary measures. The first captures *potential* exposure and is based on the occupation-level AI exposure index developed by Eloundou et al. (2024), which maps task descriptions as classified in the U.S. Department of Labor's *Occupational Information Network* (O*NET) database into predicted AI applicability using machine learning techniques and human validation. The second measure captures *realized* exposure using the so-called Anthropic Economic Index, which aggregates millions of observed interactions with *Claude* – a largely diffused LLM platform – and provides information both on the intensity of AI use by occupation and on whether enquiries to *Claude* are primarily directed to automating or augmenting labor (Handa et al., 2025).

The empirical analysis yields a rich set of findings, that BCC classifies under six main categories. First, employment among workers aged 22-25 declines significantly in occupations with high AI exposure. These declines emerge in late 2022 and are concentrated in occupations such as software development and customer service. In contrast, employment among young workers in low-exposure occupations continues to grow over the same period. Second, these occupation-level effects translate into a marked age gradient in aggregate employment outcomes. While overall employment continues to expand, employment among workers aged 22-25

[4] OpenAI has made available ChatGPT to the public in November 2022.

stagnates or declines. Decomposition by AI exposure shows that this pattern is driven almost entirely by highly exposed occupations, in which employment falls by approximately 6 percent for young workers while increasing by 6-9 percent for older workers. Employment growth in low-exposure occupations is similar across age groups. Third, the employment effects differ sharply by the nature of AI use. Occupations characterized by automating uses of AI – like computer programming or accounting – experience substantial reductions in entry-level employment, whereas occupations in which AI is primarily used in an augmentative manner – like management or repairs – do not exhibit employment declines and in some cases show increases. Fourth, the results are robust to controls for firm-level and sectoral shocks. Event-study estimates based on Poisson models with firm-by-time fixed effects indicate a decline of approximately 12 log points in employment among 22-25 year-olds in the highest AI-exposure quintile relative to less exposed occupations. No comparable effects are observed for older workers. These findings mitigate concerns that the results reflect industry-specific downturns or firm-level shocks unrelated to AI adoption. Fifth, the observed adjustment occurs along the employment margin rather than through wages. BCC finds no systematic differences in real wage growth by age or AI exposure over the sample period. This pattern is consistent with short-run wage rigidity and suggests that the initial impact of generative AI operates primarily through hiring and job availability rather than wage adjustment. Sixth, the findings are robust across a wide range of alternative specifications and samples. The results persist when excluding IT occupations or ICT-sector firms, when accounting for tele-workability, and when extending the pre-treatment period to show that no relationship between AI exposure and employment is present prior to 2022. Similar patterns are observed in occupations that do not require a college degree, indicating that labor market experience, rather than formal education alone, mediates exposure to AI-driven displacement.

The age- and experience-specific nature of the BCC findings distinguishes them from much of the earlier literature, which has largely interpreted technological change through differences in formal skills, educational attainment, or occupational task content. In particular, the routine-biased technological change literature documents employment losses concentrated in middle-skill occupations, accompanied by employment growth at the upper and lower ends of the wage distribution (Autor et al., 2003; Goos et al., 2014). By contrast, the BCC evidence identifies a gradient operating primarily across experience rather than conventional measures of skill. This pattern is difficult to reconcile within existing frameworks because entry-level and experienced workers frequently perform tasks within the same occupations and possess comparable formal qualifications. It instead points toward mechanisms operating within firms and across career stages, whereby AI automates part of the productive content of entry-level work while simultaneously augmenting the productivity of workers engaged in more judgment-intensive activities.

Although still unpublished, the BCC paper has already stimulated a broad debate regarding both the identification of its empirical results and the economic mechanisms underlying them. One line of criticism concerns timing. Iscenko and Millet (2026) note that vacancies in highly AI-exposed occupations peaked in early 2022 and began declining several months before the public release of ChatGPT, closely coinciding with the Federal Reserve's most aggressive monetary tightening cycle in decades. Since sectors such as Finance and Information are simultaneously among the most exposed to AI and the most sensitive to capital costs, they argue that the observed contraction in entry-level hiring may primarily reflect cyclical macroeconomic conditions rather than the direct consequences of AI adoption. They further contend that focusing on workers aged 22-25 mechanically amplifies the appearance of displacement, as a generalized hiring freeze

disproportionately affects the youngest cohort while older workers continue to accumulate through normal demographic progression.

In an on-line replay,[5] Brynjolfsson and co-authors challenge both arguments. They first show that available measures of occupational sensitivity to interest rates are, if anything, negatively correlated with AI exposure, suggesting that the occupations most affected by generative AI are not those most exposed to monetary tightening. Moreover, although the initial employment effects observed during 2022 and 2023 become substantially attenuated once increasingly demanding firm-by-time fixed effects are introduced, a pronounced negative relationship between AI exposure and early-career employment re-emerges in 2024 and strengthens thereafter. Their interpretation is correspondingly more nuanced than the original "Canaries" evidence: macroeconomic conditions likely contributed to the initial contraction in hiring, but they do not appear sufficient to explain the persistent divergence in employment trajectories across occupations with different levels of AI exposure.

A second line of criticism shifts attention from macroeconomic conditions to organizational change. Lambert and Schindler (2026) argue that occupations most exposed to generative AI largely coincide with those transformed by the widespread adoption of remote work during the COVID-19 pandemic – a work arrangement that, despite its partial reversal, still accounts for roughly one quarter of working days in the United States. They document a Spearman correlation of 0.77 between AI and remote-work exposure, implying that estimates based solely on AI exposure may suffer from substantial omitted-variable bias. Indeed, once both dimensions are jointly considered, the estimated AI effect often becomes statistically negligible, whereas exposure to remote work continues to account for the contraction in entry-level hiring. Their interpretation emphasizes the economics of supervision rather than automation. By raising the costs of monitoring inexperienced workers and weakening opportunities for informal mentoring, remote work disrupts the transmission of tacit knowledge and increases the relative cost of developing junior employees. Consistent with this view, Emanuel et al. (2026) show that physical proximity increases coding feedback received by junior software engineers by 18.3 percent, suggesting that firms operating remotely rely more heavily on hiring experienced workers than on cultivating talent internally.

Yet the organizational explanation should not be interpreted as an alternative to AI. Rather, it identifies the environment within which generative AI subsequently diffused. Entry-level white-collar jobs have traditionally combined two complementary functions: the production of standardized knowledge work and the accumulation of experience through apprenticeship. The widespread adoption of remote work weakened this bundle from both directions. It increased the costs of supervising and training inexperienced workers while simultaneously inducing firms to codify an expanding share of office production through digital workflows, collaborative platforms, shared documents, ticketing systems, and customer relationship management software. Although these organizational investments preceded the widespread diffusion of generative AI, they substantially expanded the stock of machine-readable activities on which AI systems could subsequently operate. Schubert (2025) provides evidence consistent with this sequencing, showing that firms investing more intensively in remote-work technologies during the pandemic were also more likely to undertake subsequent organizational investments in AI. Remote work and generative AI therefore appear as successive stages of the same organizational transformation: the former reshaped production by codifying office work, while the latter

[5] Accessible at https://digitaleconomy.stanford.edu/news/canaries-interest-rates-and-timinga-more-on-recent-drivers-of-employment-changes-for-young-workers/

exploited this newly created organizational capital. Even where firms have subsequently reduced remote work, the digital infrastructures, codified workflows, and production routines established during that period have largely remained in place, continuing to facilitate AI adoption and making this transformation largely persistent.

This organizational evolution also alters firms' incentives to invest in human capital. Garicano and Rayo (2026) argue that generative AI introduces an "AI floor" – a level of capability at which AI autonomously performs many of the routine cognitive activities historically assigned to junior employees. These activities have traditionally generated the productivity surplus financing apprenticeship. As AI progressively absorbs them, the private return to hiring and training inexperienced workers declines even though the long-run value of expertise remains unchanged. Technological adoption consequently affects labor markets through a second channel operating alongside contemporaneous labor substitution: it changes the incentives governing the production of future human capital. Consistent with this mechanism, Hosseini and Lichtinger (2025) document a pronounced divergence between junior and senior employment among firms adopting generative AI beginning in early 2023, driven almost entirely by lower hiring rather than higher separations.

All in all, these contributions reveal a remarkably consistent picture. The earliest labor-market adjustments associated with generative AI are concentrated among workers performing standardized, codifiable tasks at the beginning of their careers, although the magnitude and timing of these effects remain shaped by the broader organizational and macroeconomic environment in which AI is adopted (Greenan et al., 2025; Guarascio and Reljic, 2025). Monetary tightening, the diffusion of remote work, and the expansion of AI therefore appear as complementary forces operating at different stages of the same adjustment process. Remote work reorganized production and expanded the stock of machine-readable activities; tighter monetary conditions amplified the initial contraction in hiring; generative AI subsequently exploited this organizational environment by automating an increasing share of routine cognitive work.

The central analytical question is how AI adoption reshapes firms' incentives to reorganize production and invest in human capital, and how these decisions propagate into aggregate productivity, labor-market dynamics, and long-run growth (Ide, 2026). This perspective is closely aligned with Akin's (2026) notion of an apprenticeship externality, whereby firms' privately optimal hiring and automation decisions lead to insufficient investment in expertise from the standpoint of the economy as a whole. Once expertise is recognized as a productive input into cognitive augmentation rather than merely an attribute of individual workers, the consequences of AI extend well beyond the contemporaneous allocation of labor across tasks. They directly affect the economy's capacity to accumulate human capital, create new productive activities, and sustain productivity growth over time. The presence of an externality also establishes a natural rationale for public intervention aimed at bringing private incentives closer to the social value of expertise accumulation and organizational innovation. It is this transition – from the static allocation of tasks to the dynamic production of expertise and its macroeconomic consequences – that motivates the theoretical framework developed in the remainder of the paper.

3. Displacement, reinstatement and cognitive augmentation

The empirical regularities discussed above suggest that no existing theoretical framework, taken in isolation, is sufficient to account for the early labor-market effects of generative AI. Instead, the model developed in this paper builds on and integrates two influential, yet largely separate,

strands of the literature. The first is the automation-reinstatement (AR) approach of Acemoglu and Restrepo (2018, 2019a, 2019b, 2022), which analyzes technological change through the allocation of tasks between labor and capital. The second is the decision-theoretic perspective on AI as a form of cognitive augmentation (CA), developed by Agrawal et al. (2018, 2019a, 2019b, 2025, 2026) and Autor and Thompson (2025), which interprets AI primarily as a technology reducing the cost of prediction and extending the reach of human judgment.

Within the AR framework, technological progress operates by reallocating tasks between labor and capital. Automation reduces labor demand in existing activities, while the endogenous creation of new tasks allows labor to retain a comparative advantage in previously unavailable forms of production. Long-run employment therefore depends on the balance between technological encroachment into existing tasks and the expansion of the task frontier. The CA perspective focuses on a different dimension of technological change. By lowering the cost of prediction and inference, AI increases the productivity of workers engaged in activities requiring judgment, contextual interpretation, and organizational coordination. Technological change consequently reshapes not only the allocation of tasks but also the allocation of decision-making within firms, altering the relative value of different forms of human expertise. As emphasized by Autor and Thompson (2025), AI extends the effective reach of expertise rather than simply replacing it, allowing experienced workers to supervise a broader range of activities while codifying elements of expert reasoning that can be deployed throughout the organization.

Taken together, these perspectives describe complementary dimensions of the same technological process. Their integration yields three margins through which AI reshapes comparative advantage across the task space: automation expands the range of existing activities that can be performed by capital; cognitive augmentation raises effective labor productivity where human judgment remains valuable; and the relaxation of cognitive constraints makes new human activities economically viable, expanding the task frontier. The resulting allocation of production therefore evolves through the displacement of labor from automatable tasks, productivity gains within those that remain human-performed, and the endogenous creation of new tasks in which labor acquires or retains comparative advantage.

This synthesis also provides a coherent interpretation of the organizational adjustments documented in the recent empirical literature. Once task allocation and decision-making are considered jointly, automating and augmenting uses of AI become two dimensions of the same process. Employment declines where AI substitutes for routine cognitive activities and allows decision rights to shift from junior workers toward algorithms or more experienced employees, while productivity gains are more likely to dominate where AI complements judgment-intensive work. The muted short-run response of wages is consistent with adjustment occurring initially through hiring, workflow redesign, and changes in career structures, with compensation responding more gradually as firms reorganize production.

These adjustments acquire a dynamic dimension once entry-level work is recognized as an input into human capital formation. Junior positions have traditionally combined current production with learning by doing, with the surplus generated by routine cognitive activities helping to finance the accumulation of expertise. Automation erodes this productive component at the same time as cognitive augmentation raises the value of experienced judgment, weakening firms' incentives to hire and train inexperienced workers. The resulting wedge between the private return to apprenticeship and the social value of expertise accumulation gives rise to the apprenticeship externality introduced above. AI adoption consequently affects both current labor demand and the process through which the capabilities required for future production are formed.

The same interaction gives rise to seniority-biased technological change. By reducing demand for the activities disproportionately performed by junior workers while increasing the productivity of supervision, interpretation, and contextual judgment, AI shifts relative labor demand toward workers who already embody greater expertise. Seniority bias and the apprenticeship externality are thus closely connected: the former describes the reallocation of demand toward experienced labor, while the latter captures the dynamic consequence of that reallocation for the future supply of expertise. As training opportunities contract, expertise becomes an endogenous state variable through which current adoption decisions propagate into subsequent productivity, wage inequality, labor-market dynamics, and factor shares.

The framework developed below formalizes these relationships within a common structure, linking automation, cognitive augmentation, task creation, and expertise accumulation to the transmission of firm-level AI adoption into aggregate outcomes. Its purpose is to isolate the structural margins governing this transmission and, in doing so, provide analytical discipline for a broader research agenda on the macroeconomic consequences of AI diffusion.

3. In search of a unified framework

At the core of the AR-AC synthesis lies the idea that automation and augmentation coexist across a heterogeneous task space. In the task-based framework pioneered by Zeira (1998) in the context of capital-skill complementarity, aggregate output is the result of a continuum of discrete tasks, rather than a direct combination of capital and labor inputs, while technical change enables the displacement of labor by capital across an expanding range of tasks instead of merely augmenting factor productivity (Acemoglu and Autor, 2011). The equilibrium distribution of tasks is determined by the relative costs of factors and the endogenous boundary of the technology frontier, facilitating a rigorous treatment of both the displacement effect and the countervailing reinstatement effect through the continuous creation of new, labor-intensive tasks (Acemoglu and Restrepo, 2018). Within such a framework, we argue that AI simultaneously: (*i*) expands the set of tasks that can be performed by AI-endowed capital, displacing altogether human cognition in some domains; (*ii*) raises the effective productivity of workers in remaining tasks by reducing their cognitive constraints and/or expanding their judgment and deliberative capabilities.

Consider a competitive economy producing a single final good Y by combining a continuum of tasks indexed by $i \in [0, N]$. The final good is produced using a constant elasticity of substitution (CES) aggregator over tasks:

$$Y = \left[\int_0^N y(i)^{\frac{\sigma-1}{\sigma}} di\right]^{\frac{\sigma}{\sigma-1}} \tag{1}$$

where $\sigma > 0$ denotes the elasticity of substitution across tasks. A higher σ implies greater scope for reallocating production across tasks in response to technological change. Each task i can be produced using either labor or AI-capital. Factor markets are competitive, and firms choose the cost-minimizing method of production for each task given prevailing factor prices and technologies.

AI is modelled as a dual-use technology operating through two distinct but interacting channels. First, following the AR approach, AI acts as an automation technology that allows capital to perform an expanding subset of tasks previously carried out by labor. Second, in line with the CA approach, AI also enhances the productivity of workers in non-automated tasks by improving prediction, information processing, and decision quality. Formally, task production is given by:

$$y(i) = \begin{cases} \alpha_K(i)\, k(i), & \text{if } i \leq I \\ \alpha_L(i)\, \beta(i)\, l(i), & \text{if } i > I \end{cases} \tag{2}$$

where $k(i)$ and $l(i)$ denote capital and labor inputs, respectively. Labor is assumed to be task-specific: for each task i there exists a fixed, inelastically supplied quantity of labor $\bar{l}(i)$ that can only be employed in that task. This assumption – which can be rationalized by task-specific human capital, institutional rigidities, and frictions in reskilling and occupational mobility – ensures the existence of a non-degenerate equilibrium wage distribution, a feature that will play a central role in the analysis below.

Tasks requiring only AI-empowered capital tools are characterized by productivity $\alpha_K(i)$. For tasks performed by labor, instead, productivity depends on two components. The first term $\alpha_L(i)$ captures baseline task-specific productivity, while $\beta(i) \geq 1$ captures cognitive augmentation – the extent to which AI acts as a "bicycle for the mind" – amplifying the effective units of labor by lowering cognitive frictions. A formal derivation is provided below. Empirically, $\beta(i)$ can be proxied using measures of task uncertainty, information complexity, or reliance on judgment, while $\alpha(i)$ can be captured by experience, education, or domain expertise. The model therefore states that AI adoption affects labor productivity primarily where these dimensions interact.[6]

The threshold $I \in [0, N]$ defines the automation frontier that for given task-specific wages $W(i)$ and an interest rate exogenously fixed at R, is endogenously determined by the condition $\frac{R}{\alpha_K(I)} = \frac{W(I)}{\alpha_L(I)\beta(I)}$. Tasks with index $i \leq I$ are fully automated thanks to LLMs or ML tools, while tasks with $i > I$ remain labor-intensive endeavors. Provided that $\frac{\alpha_K(i)}{\alpha_L(i)\beta(i)}$ is decreasing in i – meaning that lower-indexed tasks are more suitable for automation – threshold I exists and is unique.

Here lies the central novelty of our framework, bridging the AR task-based approach with the AGG cognitive augmentation perspective. Once labor productivity is allowed to be enhanced by task-specific augmentation factors $\beta(i)$, comparative advantage is no longer governed by baseline human productivity $\alpha_L(i)$ alone but by effective productivity $\alpha_L(i)\beta(i)$. Crucially, baseline human productivity need not be monotone in task complexity to ensure existence and uniqueness of an automation frontier: sufficiently steep increases in augmentation can restore a well-ordered task frontier even when intrinsic skills are hump-shaped or locally declining. Cognitive augmentation reorders tasks by amplifying human advantage where judgment, interpretation, and context are most valuable. Tasks that fail to benefit from this amplification are precisely those for which augmentation grows too slowly to offset declining baseline productivity.

Finally, N denotes the measure of tasks that are technologically feasible to be performed in production at a given point in time. Tasks are indexed on a continuum, and N represents the upper bound of this task space, which expands endogenously with technological progress. An increase in N captures the arrival of new tasks – typically requiring labor input – that become economically

[6] The cognitive augmentation factor $\beta(i)$ enters the task-level production function in a manner closely related to the matching technology proposed by Althoff and Reichardt (2026). In their framework, task-level output obtained through labor depends on a matching function $m(h, r)$, where h denotes the vector of workers' skills and r the set of skill requirements needed to perform a given task. Technological progress – such as that induced by AI adoption – operates by reducing r, the level of skills required for task completion. This process, which they refer to as *simplification*, expands the set of workers who can effectively perform the task and unambiguously attenuates the returns to skill differences. The microfoundation we propose retains greater flexibility. Rather than assuming that technological change uniformly lowers task requirements, cognitive augmentation operates by scaling effective productivity within tasks in a task-specific manner. This allows augmentation to differentially affect tasks and workers, potentially compressing or amplifying wage dispersion depending on how improvements in prediction and judgment interact with baseline productivity and the structure of task substitutability.

relevant as the production frontier advances, thereby allowing for both automation of existing tasks and the creation of new ones within the economy.

As discussed above, the central empirical finding of BCC is that AI exposure does not map cleanly onto skill or occupation categories. Instead, effects vary sharply across tasks and across workers within the same occupation. This evidence can be rationalized inside this framework by distinguishing three channels through which AI impacts on the labor market and aggregate outcome, the first and third ones ensuing from the AR approach and the second one from the AGG approach:

1. **Displacement** ($\Delta I > 0$). When AI expands the automation frontier without significantly augmenting remaining tasks, labor demand falls in routine cognitive activities. Wages and employment decline for tasks characterized by standardization and codifiability, such as entry-level writing, translation, and basic coding.
2. **Cognitive Augmentation** ($\Delta\beta(i) > 0$). For tasks that remain labor-intensive, AI raises the marginal product of labor by improving prediction quality and reducing errors. BCC show that productivity gains are often largest for lower- and mid-performing workers – the "canaries in the coal mine" – whose output quality converges toward that of top performers when supported by AI tools.
3. **Reinstatement** ($\Delta N > 0$). AI also generates new tasks that require human judgment, domain knowledge, and coordination. Examples include roles such as AI orchestrators, prompt engineers, and compliance specialists. These tasks reinstate labor demand at the upper end of the task distribution and can offset displacement effects elsewhere.
   We now proceed to formally derive the augmentation factor $\beta(i)$ from first principles, and to relate it to the notion of expertise.

3.1. Tasks as prediction-dependent decision problems

The reduced-form representation introduced in the previous section captures a fundamental insight of the microeconomics of AI: modern AI systems primarily improve human-based activities and decision-making by improving *prediction*. The CA approach formalizes this insight by showing that a fall in the cost of prediction changes the optimal organization of decision-making and raises productivity in tasks that rely on inference under uncertainty. Here we show how this framework provides a natural microfoundation for the task-specific augmentation factor $\beta(i)$, and how $\beta(i)$ interacts with baseline human productivity $\alpha(i)$ to shape comparative advantage and task allocation.

Consider task $i \in (I,\ N)$ as a decision problem in which a human worker must choose an action $a \in A$ based on an uncertain state $\theta \in \Theta$. Per-worker output in task i is produced according to:

$$q(i) = \alpha_L(i) f(a, \theta) \tag{3}$$

where $\alpha_L(i)$ captures baseline human skill – experience, domain knowledge and contextual understanding – while $f(\cdot)$ represents *judgment*, i.e. the activity of mapping actions and realized states into output. Because the worker does not observe θ directly, expected output conditional on information is:

$$E[q(i)] = \alpha_L(i) E[f(a, \theta)]. \tag{4}$$

The quality of the chosen action depends on the accuracy of the worker's prediction $\hat{\theta}$ of the true state θ. Following Agrawal et al. (2018, 2025), the key input into the decision is therefore a prediction $\hat{\theta}$ about the underlying state. Prediction is costly and imperfect, meaning that it generates errors that reduce expected output. Let the expected loss from prediction error be:

$$\mathcal{L}_i = \mathrm{E}[(\theta - \hat{\theta})^2], \tag{5}$$

that depends on the availability and quality of AI-based prediction tools.

Different tasks rely on prediction to different degrees. Let $\lambda(i) > 0$ denote the prediction intensity of task i, defined as the marginal sensitivity of output to prediction errors. Prediction errors degrade judgment, and degradation is task-specific: tasks involving diagnosis, forecasting, prioritization, or coordination under uncertainty have high $\lambda(i)$, while routine tasks have low $\lambda(i)$. We assume that expected judgment can be approximated by a first-order loss expansion around the optimum, so that we can write:

$$E[f(a,\theta)] = \bar{f}(i) - \lambda(i)\mathcal{L}_i \tag{6}$$

where $\bar{f}(i)$ is the output that would be achieved under perfect prediction.

Improvements in AI technology reduce the marginal cost of prediction, thereby lowering $\mathcal{L}_i$. Let $\mathcal{L}_i^0$denote the prediction error without AI and $\mathcal{L}_i$the error with AI assistance, where $\mathcal{L}_i < \mathcal{L}_i^0$. Using (4), the expected value of per-worker output in the two cases are:

$$Without\ AI \qquad E[q^0(i)] = \alpha_L(i)\left[\bar{f}(i) - \lambda(i)\mathcal{L}_i^0\right] \tag{7}$$

$$With\ AI \qquad E[q(i)] = \alpha_L\left[\bar{f}(i) - \lambda(i)\mathcal{L}_i\right]. \tag{8}$$

By combining (7) and (8), after some simple algebra, we obtain:

$$E[q(i)] = E[q_0(i)]\left[1 + \frac{\lambda(i)}{\bar{f}(i) - \lambda(i)\mathcal{L}_i^0}(\mathcal{L}_i^0 - \mathcal{L}_i)\right]. \tag{9}$$

Expression (9) makes clear that AI acts as a *multiplicative shifter* of baseline expected output. This motivates defining $\beta(i)$ accordingly as:

$$\beta(i) \equiv 1 + \frac{\lambda(i)}{\bar{f}(i) - \lambda(i)\mathcal{L}_i^0}(\mathcal{L}_i^0 - \mathcal{L}_i). \tag{10}$$

Let $D_i = \bar{f}(i) - \lambda(i)\mathcal{L}_i^0 > 0$, so that the expected output in the no-AI benchmark is assumed to be positive. Holding $\lambda(i)$ fixed, the sensitivity of $\beta(i)$ to $\mathcal{L}_i$ is given by:

$$\frac{\partial\beta(i)}{\partial\mathcal{L}_i} = -\frac{\lambda(i)}{D_i} < 0. \tag{11}$$

Thus, any improvement in AI prediction accuracy modeled as a reduction in $\mathcal{L}_i$ raises $\beta(i)$ linearly, with the magnitude of the effect proportional to $\lambda(i)$. This is precisely the CA prediction channel: AI acts as a general-purpose prediction technology, in that improvements in algorithms, data, or compute reduce $\mathcal{L}_i$ across many tasks. Importantly, this channel operates even when the structure of tasks and human judgment remain unchanged. AI enhances labor by narrowing the informational wedge between the chosen action and the optimal one.

In turn, holding $\mathcal{L}_i$ fixed allows us to define the sensitivity of cognitive augmentation to $\lambda(i)$ as:

$$\frac{\partial \beta(i)}{\partial \lambda(i)} = \frac{\bar{f}(i)(\mathcal{L}_i^0 - \mathcal{L}_i)}{\left[\bar{f}(i) - \lambda(i)\mathcal{L}_i^0\right]^2} > 0. \tag{12}$$

so that an increase in $\lambda(i)$ – i.e., a greater sensitivity of output to prediction errors – increases effective productivity for any given level of prediction accuracy. This static effect contains a deeper implication once AI adoption is allowed to respond endogenously. Indeed, it is immediate to note that tasks with high $\lambda(i)$ are precisely those where reductions in $\mathcal{L}_i$ generate large gains:

$$\frac{\partial^2 \beta(i)}{\partial \mathcal{L}_i \partial \lambda(i)} = -\frac{\bar{f}(i)}{\left[\bar{f}(i) - \lambda(i)\mathcal{L}_i^0\right]^2} < 0. \tag{13}$$

Thus, tasks with higher $\lambda(i)$ benefit *disproportionately* from any given improvement in AI prediction. This maps directly into the expertise approach of Autor and Thompson (2025). High-$\lambda(i)$ tasks are in general expert-intensive activities, like for instance diagnosis, evaluation, synthesis, and exception handling, where small errors have large consequences.[7] In such tasks, expertise is scarce because judgment is difficult to scale. In all these cases AI does not eliminate judgment, but by improving prediction it relaxes the constraint that makes expertise valuable and rare. This means that $\beta(i)$ is *more elastic* with respect to AI progress in tasks with high $\lambda(i)$, and expertise-intensive tasks are therefore natural early beneficiaries of cognitive augmentation.

Using (9), (7) and (10), expected per-person output becomes:

$$E[q(i)] = \alpha_L(i)\beta(i)\tilde{f}(i) \tag{14}$$

where $\tilde{f}(i) = \left[\bar{f}(i) - \lambda(i)\mathcal{L}_i^0\right]$. We normalize $\tilde{f}(i) = 1$ so that, without loss of generality by choice of units, it is absorbed into $\alpha_L(i)$. Finally, as we assume that task i employs $l(i)$ identical specialized workers, each producing expected output $\alpha_L(i)\beta(i)$, with constant returns to labor within the task one obtains:

$$y(i) = \int_0^{l(i)} \mathrm{E}[q(i)]\, dj = \alpha_L(i)\beta(i)l(i). \tag{15}$$

Equations (14)-(15) provide the precise analytical bridge between micro-level improvements in decision quality emphasized in the CA approach and the reduced-form AR task-based production function used in macroeconomic analysis of technological change. Starting from a primitive decision problem in which output depends on judgment under uncertainty, improvements in prediction ensured by AI tools translate into higher expected output through a reduction in task-specific prediction losses. This operates by shrinking the gap between potential and realized productivity, with the strength of the effect governed by the task's sensitivity to misprediction $\lambda(i)$, and by the magnitude of the AI-induced reduction in loss. The resulting augmentation factor $\beta(i)$ captures how advances in prediction technology amplify effective labor input in judgment-intensive tasks.

[7] More on this below.

3.2. Expertise, apprenticeship and seniority-biased technological change

The framework developed thus far is static, as the augmentation factor $\beta(i)$ is taken as given once its AI-related content has been determined. Yet the evidence reviewed above suggests that one of the defining features of AI concerns its complex interaction with the level and accumulation of expertise itself. Entry-level tasks contribute not only to current output, but also to the accumulation of judgment competence, contextual understanding, and tacit knowledge through which workers progressively acquire expertise (Autor and Thompson, 2025). To easily capture this dynamic dimension while preserving the structure of the model, we allow cognitive augmentation to depend on the economy's stock of accumulated expertise. By generalizing the baseline model, this simple extension delineates a wide-ranging framework from which a series of critical corollaries may logically follow.

Let E_t denote the aggregate stock of productive expertise at date t. Unlike formal education or general human capital, expertise summarizes the ability to transform AI-generated predictions into economically valuable decisions through contextual interpretation, supervision, and exception handling, and it enhances the effectiveness of AI assistance. Accordingly,

$$\beta(i,t) = \beta(\mathcal{L}_i, \lambda(i), E_t), \qquad \frac{\partial \beta(i,t)}{\partial E_t} > 0, \tag{16}$$

so that the production technology in equation (15) becomes $y(i,t) = \alpha_L(i)\beta(i,t)l(i,t)$. Equation (16) captures the complementarity between AI-generated prediction, human-based expertise and judgment: improvements in AI reduce prediction losses through $\mathcal{L}_i$, while accumulated expertise determines the extent to which these improvements translate into higher effective labor productivity associated to better decisions.

Expertise is accumulated through apprenticeship within firms: by performing tasks under the supervision of more experienced workers, junior employees acquire the tacit knowledge, contextual judgment, and organizational know-how that cannot be fully transmitted through formal education or codified instruction. Let $\Im_t \subset (I_t, N_t]$ denote the subset of non-automated tasks performed by inexperienced workers under the supervision of senior employees. These tasks jointly generate contemporaneous output and learning-by-doing, so that the aggregate stock of expertise evolves according to:

$$E_{t+1} = (1 - \delta_E)E_t + \int_{\Im} \phi(i)l(i,t)di \tag{17}$$

where $\delta_E \in (0,1)$ is the depreciation rate of expertise and $\phi(i) > 0$ measures the learning content embodied in task i. Tasks characterized by intensive interaction with experienced workers, repeated exposure to uncertainty, and progressive assumption of decision-making responsibilities therefore contribute more strongly to future expertise accumulation. An outward shift of the automation frontier reduces the measure of apprenticeship tasks contained in $\Im_t$, lowering the future stock of expertise. At the same time, a lower E_t weakens subsequent cognitive augmentation through equation (16), reducing the productivity gains that AI can generate in judgment-intensive activities. Current automation decisions therefore influence not only contemporaneous task allocation but also the future effectiveness of AI as a complement to human labor.

This dynamic naturally gives rise to the growth-constraining apprenticeship externality emphasized by Ide (2026) and Akin (2026). Junior employment jointly produces current output and expertise, but firms appropriate only part of the return to the latter. Workers carry the

knowledge and experience acquired through apprenticeship across jobs, allowing part of the resulting productivity gains to accrue to future employers and, more broadly, to the economy. Firms consequently evaluate apprenticeship against a private return that falls short of its social value, creating a wedge between privately optimal automation decisions and the level of expertise accumulation that is efficient from an aggregate perspective.

To formalize this wedge, let $V_{E,t}^{S} > 0$ denote the social marginal continuation value of one additional unit of expertise accumulated at date t. We assume that the training firm internalizes only a fraction $\chi \in [0,1]$ of this return: worker mobility, incomplete contracting, and knowledge spillovers across employers reduce χ, whereas longer employment relationships or institutional arrangements allowing firms to retain trained workers increase it. The private marginal continuation value of expertise is therefore simply given by:

$$V_{E,t}^{P} = \chi V_{E,t}^{S}. \tag{18}$$

Let $\Delta\pi A(i,t) > 0$ denote the contemporaneous surplus generated, per junior worker, by automating apprenticeship task *i*. According to the law of motion (17), retaining the task in human production instead generates $\phi(i)$ units of additional expertise, each valued privately at $V_{E,t}^{P}$. A profit-maximizing firm therefore automates whenever:

$$\Delta\pi A(i,t) > V_{E,t}^{P}\ \phi(i) = \chi V_{E,t}^{S}\phi(i), \tag{19}$$

while a social planner would internalize the full continuation value of the expertise generated through apprenticeship by choosing:

$$\Delta\pi A(i,t) > V_{E,t}^{S}\phi(i). \tag{20}$$

Clearly, whenever $\chi < 1$ the private and social automation thresholds diverge, and the automation of certain tasks generates an economy-wide externality. This is due to the fact that along a competitive equilibrium all apprenticeship-intensive tasks satisfying $\chi V_{E,t}^{S}\ \phi(i) < \Delta\pi_{A}$ $(i,t) < V_{E,t}^{S} S\phi(i)$ are automated, even though preserving them would be socially efficient. This distortion arises because firms compare the immediate productivity gains from automation with only the privately appropriable component of future expertise, whereas society values the entire stream of future productivity generated by the learning embodied in apprenticeship. Consequently, as the automation frontier expands, the economy undergoes not only a contemporaneous substitution of capital for labor across tasks but also an inefficient reduction in the accumulation of expertise. Since expertise enters the cognitive augmentation factor through equation (16), lower investment in apprenticeship today reduces the capacity of future workers to transform AI-generated predictions into productive decisions, thereby weakening the long-run complementarity between AI and human judgment. More broadly, this mechanism suggests that the efficiency of AI adoption cannot be evaluated solely in terms of its immediate productivity gains. It also depends on its dynamic effects on the economy's capacity to accumulate expertise, making the interaction between automation, human capital formation, and organizational learning a central object for future theoretical and empirical research on the long-run growth consequences of AI.

The same extension also provides a formal rationale for the emergence of seniority-biased technological change. Consider an increase in AI capability that simultaneously expands the

automation frontier ($\Delta I_t > 0$) and improves prediction quality ($\Delta \mathcal{L}_i > 0$). For every non-automated task,

$$d[\alpha_L(i)\beta(i,t)] = \alpha_L(i)\left(\frac{\partial \beta}{\partial \mathcal{L}_i} d\mathcal{L}_i + \frac{\partial \beta}{\partial E_t} dE_t\right), \tag{21}$$

where $(\partial \beta / \partial L_i) < 0$ implies that improvements in prediction raise effective labor productivity. The magnitude of this gain, however, depends on the stock of accumulated expertise. We assume that expertise and prediction are technological complements,

$$\frac{\partial^2 \beta(i,t)}{\partial E_t \partial(-\mathcal{L}_i)} > 0, \tag{22}$$

so that workers endowed with greater expertise derive larger productivity gains from a given improvement in AI prediction. This complementarity captures the idea that AI supplies increasingly accurate predictions, whereas expertise determines the capacity to interpret, validate, and translate those predictions into productive decisions.

The implications for labor demand follow immediately. The expansion of the automation frontier disproportionately substitutes for standardized tasks that constitute the productive core of apprenticeship, thereby reducing demand for inexperienced workers and lowering the future accumulation of expertise. At the same time, equation (22) implies that the remaining judgment-intensive tasks become relatively more productive precisely because they are performed by workers embodying higher levels of accumulated expertise. AI adoption therefore simultaneously contracts the lower end of the career ladder while increasing the relative productivity and demand for experienced workers. Seniority-biased technological change emerges as the equilibrium consequence of the interaction between task substitution and cognitive augmentation, while the apprenticeship externality represents its dynamic counterpart: by weakening the mechanism through which expertise is accumulated, excessive automation today reduces the economy's future capacity to exploit AI as a complement to human judgment.

4. Wage inequality

The evolution of wage distribution is one of the principal objects through which macroeconomists study the aggregate consequences of technological change (Acemoglu and Restrepo, 2022; Hémous and Olsen, 2022). Wage inequality summarizes how productivity gains are transmitted across workers, tasks, and sectors, providing a direct link between microeconomic technological change and aggregate outcomes such as labor income, structural transformation, and long-run growth. Within the framework developed here, the distribution of wages emerges endogenously from the interaction between baseline task productivity, cognitive augmentation, and the accumulation of expertise.

Recall that for tasks remaining in the human domain, production is given by:

$$y(i) = \alpha_L(i)\beta(i,t)l(i). \tag{23}$$

Under perfect competition, the demand for task *i* implied by the CES technology satisfies:

$$y(i) = Y\left(\frac{p(i)}{P}\right)^{-\sigma}, \tag{24}$$

with $p(i) = W(i)/(\alpha_L(i)\beta(i,t))$.

Unlike the standard AR task-based model, the supply of labor to each task is no longer interpreted as an exogenous endowment. Let denote the effective supply of workers capable of performing task *i*. In the short run, this supply remains locally inelastic because of occupational specialization and adjustment frictions. Over longer horizons, however, it evolves through the accumulation of expertise described in section 3.2. Formally,

$$\bar{l}(i,t) = \ell(i)Et, \qquad \ell(i) > 0, \tag{25}$$

where $\ell(i)$ measures the intensity with which aggregate expertise is embodied in task i. Tasks requiring greater judgment, contextual reasoning, and supervisory capabilities are therefore characterized by larger values of $\ell(i)$, implying that the effective supply of labor available for those activities depends directly on the economy's accumulated stock of expertise.

Normalizing the aggregate price index to unity, equilibrium wages become:

$$W(i) = (\alpha_L(i)\beta(i,t))^{\frac{\sigma-1}{\sigma}} \left(\frac{Y}{\ell(i)E_t}\right)^{\frac{1}{\sigma}}. \tag{26}$$

Equation (26) highlights the two channels through which AI affects wages. The first is contemporaneous: cognitive augmentation raises effective labor productivity through $\beta(i,t)$. The second is dynamic: the supply of expertise embodied in workers evolves through E_t, itself determined by apprenticeship. Consequently, wage determination depends simultaneously on technological progress and on the economy's ability to reproduce the expertise required to exploit that technology.

Taking logarithms:

$$\log W(i) = \text{const} + \frac{\sigma-1}{\sigma}[\log \alpha_L(i) + \log \beta(i,t)] - \frac{1}{\sigma}[\log \ell(i) + \log E_t] \tag{27}$$

so that the cross-task cross-task wage dispersion becomes:

$$Var[\log\ W(i)] = \left(\frac{\sigma-1}{\sigma}\right)^2 Var[\log\alpha_L(i) + \log\beta(i,t)] + \frac{1}{\sigma^2} Var[\log\ell(i)] -$$

$$- \frac{2(\sigma-1)}{\sigma^2} Cov[\log(\alpha L(i)\beta(i,t)), \log\ell(i)]. \tag{28}$$

Wage inequality is therefore determined not only by heterogeneity in productivity but also by the distribution of expertise across tasks. Cognitive augmentation shapes the first component by altering effective labor productivity, whereas apprenticeship governs the second by determining the long-run supply of workers capable of performing judgment-intensive activities.

The covariance term in equation (28) plays a particularly important role. When AI-induced augmentation is concentrated in tasks already relying heavily on accumulated expertise, productivity gains reinforce existing comparative advantages while the supply of suitably trained workers adjusts only gradually. Wage dispersion consequently increases because productivity growth and expertise scarcity move in the same direction. Conversely, when cognitive augmentation disproportionately benefits tasks requiring relatively little accumulated expertise, productivity gains become more broadly distributed across the task space, attenuating existing

wage differentials. Long-run inequality therefore depends not simply on the diffusion of AI, but on how cognitive augmentation co-varies with the distribution of expertise across production.

The elasticity of substitution across tasks determines the extent to which these productivity differences translate into equilibrium wage differentials. When tasks are highly substitutable ($\sigma > 1$), firms rapidly reallocate production toward activities experiencing the largest productivity gains, amplifying wage dispersion generated by differences in effective productivity and expertise. When tasks are gross complements ($0 < \sigma < 1$), productivity improvements diffuse throughout the production network, reducing the sensitivity of equilibrium wages to local productivity shocks. The same technological improvement may therefore generate markedly different distributional outcomes depending on the organization of production.

This framework also provides a natural interpretation of seniority-biased technological change. Workers with greater accumulated expertise benefit disproportionately from cognitive augmentation because expertise increases the effectiveness with which AI-generated predictions are translated into productive judgment. At the same time, automation reduces employment precisely in the apprenticeship-intensive tasks through which future expertise is accumulated. AI therefore simultaneously raises the productivity of experienced workers while reducing the flow of workers acquiring expertise. The resulting wage premium for senior workers is not imposed exogenously but emerges endogenously from the interaction between cognitive augmentation and the evolving supply of expertise embodied in the workforce.

The apprenticeship externality amplifies this mechanism. As shown in section 3.2, firms capture only part of the future return generated by apprenticeship and therefore automate learning-intensive tasks more aggressively than is socially efficient. Lower apprenticeship reduces the future stock of expertise E_t, thereby decreasing the effective supply of workers capable of performing judgment-intensive activities. Through equation (26), this reduction raises the equilibrium wage of the remaining experts while simultaneously slowing the growth of cognitive augmentation in subsequent periods. The economy therefore experiences a dynamic feedback: insufficient investment in expertise not only limits future productivity growth but also increases the scarcity premium earned by experienced workers. Long-run wage inequality becomes the equilibrium outcome of two mutually reinforcing forces: AI increases the relative productivity of expertise, while the apprenticeship externality reduces its future supply.

Viewed from this perspective, seniority-biased technological change and the apprenticeship externality emerge as two closely connected manifestations of the same dynamic process. The former captures the contemporaneous reallocation of labor demand toward workers embodying greater expertise, while the latter describes how this reallocation weakens the incentives through which expertise is reproduced over time. Their interaction makes the distributional consequences of AI depend on the pace of automation and cognitive augmentation together with the institutional and organizational arrangements governing expertise accumulation. Understanding the long-run evolution of wage inequality therefore requires jointly identifying the dynamics of cognitive augmentation, the elasticity of substitution across tasks, and the economy's capacity to sustain the stock of expertise on which the productivity of human–AI complementarity increasingly depends.

5. Factor shares and balanced growth

The distribution of wages is only one dimension through which macroeconomists evaluate the aggregate consequences of technological change. An equally important question concerns the evolution of factor shares, since the division of income between labor and capital provides one of

the principal benchmarks against which theories of long-run growth are assessed. Since Kaldor's (1961) stylized facts, explaining how sustained productivity growth can coexist with a persistent role for labor in production has remained a defining challenge for growth theory. Recent evidence documenting a secular decline in labor's share has renewed this question, raising the possibility that automation and artificial intelligence may fundamentally alter the mechanisms through which technological progress propagates into the aggregate economy (Autor et al., 2020; Grossman and Oberfield, 2022). Within the standard neoclassical framework, balanced growth requires technical change to be purely labor augmenting. Task-based models relax this restriction by allowing automation and the endogenous creation of new tasks to jointly determine the evolution of labor demand. The AR–CA framework developed here introduces a further margin of adjustment through cognitive augmentation while recognizing that the long-run productivity of labor depends on the endogenous accumulation of expertise. The conditions under which AI remains compatible with balanced growth therefore depend on the joint evolution of automation, cognitive augmentation, task creation, and expertise accumulation.

To characterize the asymptotic properties of the model, consider an equilibrium in which aggregate output, the task frontier, and effective labor productivity all grow at constant rates, while the fraction of automated tasks converges to a constant. Let I_t denote the automation frontier and N_t the upper bound of the task frontier. Balanced growth therefore requires the ratio $I_t \,/\, N_t$ to be asymptotically stationary.

To characterize the long-run evolution of the automation frontier, it is useful to summarize the productive capability of AI-enabled capital through a single object. Let

$$\gamma(i,t) = \frac{\alpha_K(i,t)}{R_t}$$

be the measure of the effective automation capability of AI in task i, namely the quantity of task output that AI-enabled capital can produce per unit of financing cost. Unlike $\beta(i,t)$, which captures the productivity gains generated when AI complements human labor, $\gamma(i,t)$ summarizes AI's capacity to substitute for labor economically. Improvements in $\gamma(i,t)$ may originate from advances in AI models, reductions in computing costs, cheaper data infrastructure, or lower financing costs associated with AI-capital. The relevant margin for automation is therefore not technological capability alone, but its productivity relative to cost. The marginal task satisfies:

$$\gamma(I_i,t) = \frac{\alpha_L(I_t)\beta(I_t,t)}{W(I_t,t)}, \tag{29}$$

meaning that automation occurs until the output produced by one euro (or dollar) invested in AI exceeds the output produced by one euro (or dollar) spent on effective labor. Equation (29) immediately shows that automation depends on the joint evolution of AI capability, cognitive augmentation, and equilibrium wages. Improvements in AI capability alone do not determine the expansion of automation; what matters is their evolution relative to the effective productivity of augmented labor.

For the share of automated tasks to remain asymptotically constant, the automation frontier and the task frontier must expand proportionally,

$$\frac{\dot{I}_t}{I_t} = \frac{\dot{N}_t}{N_t}. \tag{30}$$

This condition guarantees that automation continues without progressively exhausting the set of tasks performed by labor. To characterize the evolution of labor's income share, recall that aggregate labor income is obtained by integrating wages over all labor-performed tasks,

$$W_t L_t = \int_{I_t}^{N_t} W(i,t)\bar{l}(i,t)di \tag{31}$$

In section 3.2 the effective labor supply $\bar{l}(i,t)$ was linked to the stock of accumulated expertise. For the purpose of the balanced-growth analysis, however, that cross-task heterogeneity plays no independent role. We therefore adopt the simplifying assumption that expertise scales the effective labor supply proportionally across all human-performed tasks, allowing its aggregate effect to be summarized by the common state variable E_t. This normalization preserves the growth implications of expertise accumulation while avoiding the introduction of additional distributional terms that are not essential for the asymptotic analysis.

Substituting the equilibrium wage schedule into (31) then yields:

$$W_t L_t = Y_y^{\frac{1}{\sigma}} E_t^{\frac{\sigma-1}{\sigma}} \int_{I_t}^{N_t} [\alpha_L(i)\beta(i,t)]^{\frac{\sigma-1}{\sigma}} di, \tag{32}$$

and dividing by aggregate output gives the labor share:

$$s_L(t) = Y_t^{-\frac{\sigma-1}{\sigma}} E_t^{\frac{\sigma-1}{\sigma}} \int_{I_t}^{N_t} [\alpha_L(i)\beta(i,t)]^{\frac{\sigma-1}{\sigma}} di. \tag{33}$$

Equation (33) highlights the three mechanisms governing long-run factor shares. Automation changes the measure of labor-performed tasks through the evolution of I_t; cognitive augmentation raises effective labor productivity through $\beta(i,t)$; and expertise accumulation shifts the effective supply of labor through E_t. Stable factor shares therefore require these three margins to evolve in a coordinated manner.

Differentiating (33) with respect to time yields:

$$\frac{\dot{s}_L}{s_L} = \frac{\dot{M}_t}{M_t} + \frac{\sigma-1}{\sigma}(g_E - g) \tag{34}$$

where $M_t = \int_{I_t}^{N_t} [\alpha_L(i)\beta(i,t)]^{\frac{\sigma-1}{\sigma}} di$, and g_E and g_Y denote the growth rates of expertise and aggregate output, respectively. Differentiating over time while accounting for the movement of both task boundaries and changes in productivity within the surviving labor-performed tasks yields:

$$\frac{\dot{M}_t}{M_t} = \frac{m(N_t,t)\dot{N}_t - m(I_t,t)\dot{I}_t}{M_t} + \frac{\sigma-1}{\sigma}\bar{g}_\beta \tag{35}$$

where $m(i,t) = [\alpha_L(i)\beta(i,t)]^{\frac{\sigma-1}{\sigma}}$, and $\bar{g}_\beta$ denotes the weighted average growth rate of cognitive augmentation across labor-performed tasks.

Combining (34) and (35), the labor share remains asymptotically constant if and only if:

$$\frac{m(N_t,t)\dot{N}_t - m(I_t,t)\dot{I}_t}{M_t} + \frac{\sigma-1}{\sigma}\left(\bar{g}_\beta + g_E - g_Y\right) = 0. \tag{36}$$

This decomposition has a useful economic interpretation. The first term is the net task-reinstatement component: $m(N_t, t)\dot{N}_t$ measures the contribution of newly created human tasks, while $m(I_t, t)\dot{I}_t$ measures the labor-income content lost through automation. The second term captures the intensive-margin contribution of human productive capacity: cognitive augmentation raises productivity within surviving tasks through $\bar{g}_\beta$, while expertise accumulation expands effective labor supply through g_E; both must be evaluated relative to aggregate output growth g_Y. Equation (36) therefore summarizes the balanced-growth condition implied by the AR–CA framework. Stable factor shares require the expansion of the task frontier to offset the progressive automation of existing tasks while the combined growth of cognitive augmentation and expertise accumulation keeps pace with aggregate productivity growth. Unlike the canonical Uzawa condition, balanced growth therefore does not rely exclusively on labor-augmenting technical change. Nor does it depend solely on the endogenous creation of new tasks emphasized by the automation-reinstatement literature. Instead, it emerges from the joint evolution of automation, cognitive augmentation, expertise accumulation, and task creation.

The apprenticeship externality gives this condition an important dynamic interpretation. When firms underinvest in apprenticeship, expertise accumulates more slowly than is socially efficient. Lower expertise weakens subsequent cognitive augmentation and reduces the economy's capacity to generate new judgment-intensive tasks, slowing the expansion of the task frontier. Consequently, AI adoption may remain privately optimal while failing to satisfy the balanced-growth condition because the economy systematically underproduces the expertise required to sustain the complementarity between AI and human labor.

The model therefore admits transitional declines in the labor share together with asymptotic stability. During the diffusion of AI, automation may initially outpace both expertise accumulation and task creation, causing labor's share to fall despite sustained productivity growth. Over time, however, continued cognitive augmentation, the endogenous expansion of the task frontier, and the accumulation of expertise restore the productive role of labor. Balanced growth is thus achieved not because automation slows, but because the economy eventually generates sufficient augmentation, expertise, and new human tasks to offset the displacement associated with technological progress.

Equation (36) also provides a useful perspective through which to interpret recent proposals by Acemoglu et al. (2026) advocating a transition toward *pro-worker* AI. Their argument starts from the observation that current patterns of AI development are disproportionately oriented toward automation, reflecting the incentives faced by firms rather than the direction of technological change that maximizes long-run social welfare. AI applications complementing workers' capabilities, facilitating the diffusion of expertise, and supporting the emergence of new human tasks receive comparatively less investment, despite their potentially larger contribution to productivity and shared prosperity.

The framework developed here gives this argument a precise macroeconomic interpretation. The evolution of automation, cognitive augmentation, expertise accumulation, and task creation jointly determines whether AI remains compatible with stable factor shares over time. Left to decentralized market forces, however, there is no reason to expect these margins to evolve in a coordinated manner. The apprenticeship externality identified above implies that firms systematically underinvest in the production of expertise because they capture only a fraction of its future return. Since expertise simultaneously increases the productivity of AI-assisted labor and supports the emergence of new judgment-intensive tasks, the resulting equilibrium tends to generate too little augmentation and too little task creation relative to the socially efficient

allocation. Automation remains privately profitable throughout this process, yet the aggregate economy gradually loses part of the productive capacity required to sustain balanced growth.

This interpretation naturally identifies the margins through which policy can improve long-run outcomes. Measures affecting the composition of innovation influence the balance between automation-oriented and augmentation-oriented AI. Investments in education, apprenticeship, continuous training, and organizational learning strengthen the accumulation of expertise, increasing the productivity of AI-assisted labor while expanding the future supply of workers capable of performing judgment-intensive activities. Policies facilitating experimentation with new organizational forms, entrepreneurial entry, and complementary digital infrastructure accelerate the emergence of new tasks, reinforcing the extensive margin through which labor remains an active factor of production as AI diffuses throughout the economy.

These interventions influence the same set of mechanisms governing the asymptotic evolution of the economy. Their common purpose is to strengthen the complementarities between automation, cognitive augmentation, expertise, and task creation, allowing productivity gains generated by AI to propagate without progressively reducing labor's contribution to production. Within this perspective, pro-worker AI is less a constraint on technological progress than a characterization of its direction. The central objective is to align private incentives more closely with the social value of expertise accumulation and organizational innovation, thereby reducing the wedge introduced by the apprenticeship externality and improving the conditions under which the economy converges toward a balanced-growth path.

6. Disciplining the empirical analysis of AI

The theoretical analysis developed so far provides useful discipline for empirical work by isolating the economic margins and structural parameters governing the transmission of technological change from individual production decisions to aggregate outcomes. Existing empirical research on generative AI has naturally begun with measures of occupational and task exposure, which have proved invaluable for mapping the reach of the technology and documenting heterogeneous responses across firms, occupations, and workers. Yet exposure remains a reduced-form object: it indicates where AI is likely to matter without revealing how adoption translates into productivity, labor demand, wage inequality, or long-run growth. The framework developed here bridges this gap by tracing these outcomes to a parsimonious set of structural parameters governing automation, cognitive augmentation, task creation, and expertise accumulation, thereby identifying the objects whose estimation is required to understand how firm-level AI adoption propagates through the aggregate economy.

The first object of interest is the cognitive augmentation factor $\beta(i,t)$, which captures the increase in effective labor productivity generated by AI through improvements in prediction and information processing. Within the AR–CA framework, $\beta(i,t)$ is an equilibrium outcome shaped jointly by AI capabilities, the prediction intensity of the task, and the expertise workers bring to the production process. Similar levels of AI exposure may therefore generate markedly different productivity gains across occupations when prediction plays a different role in production or workers differ in their ability to translate AI-generated predictions into productive judgment. By the same logic, relatively less exposed occupations may experience substantial gains where AI strongly complements judgment-intensive activities. The relevant empirical object is thus the extent to which AI adoption translates into effective labor productivity, rather than technological

exposure alone. A first-order approximation to equilibrium labor demand makes this distinction explicit:

$$\Delta \log l(i,t) = \theta_1 \Delta \log \beta(i,t) - \theta_2 \Delta \log W(i,t) + \varepsilon_{it}, \qquad \theta_1, \theta_2 > 0, \tag{37}$$

where θ_1 measures the elasticity of labor demand with respect to cognitive augmentation and θ_2 captures the offsetting effect of endogenous wage adjustment. Employment therefore responds to AI only indirectly, through changes in effective labor productivity. Estimating θ_1 requires measures of realized augmentation rather than technological availability, implying that indicators based on observed AI usage, task-level deployment, or firm-level adoption are more informative than exposure measures constructed exclusively from occupational characteristics.

The model also identifies a second empirical object closely related to cognitive augmentation. Equations (10)–(16) imply that identical improvements in AI prediction need not generate identical productivity gains across tasks. The relevant source of heterogeneity is the prediction-intensity parameter $\lambda(i)$, which measures the extent to which prediction errors constrain performance in each activity. Holding expertise constant, higher values of $\lambda(i)$ imply larger improvements in effective labor productivity following a given reduction in prediction loss. Estimating the joint distribution of $\beta(i,t)$, $\lambda(i)$, and accumulated expertise therefore becomes the first empirical step in distinguishing tasks in which AI primarily complements labor from those in which it merely substitutes for routine production.

Once these local productivity effects have been identified, the remaining question concerns their transmission to the aggregate economy. The same increase in cognitive augmentation may generate very different employment and wage responses depending on how production is organized across tasks. Capturing this propagation mechanism requires identifying the elasticity of substitution across tasks σ. While cognitive augmentation determines how AI changes effective labor productivity within a given task, σ governs how these local productivity gains propagate throughout the production system. It therefore determines the extent to which technological change is transmitted into labor demand, wage dispersion, and aggregate productivity. In this sense, σ represents the principal macroeconomic transmission mechanism of the AR–CA framework.

Its empirical importance follows directly from the equilibrium wage distribution. As shown in equation (28), wage dispersion depends jointly on the distribution of effective productivity and on the elasticity with which firms reallocate production across tasks. Estimating σ is therefore necessary to distinguish two conceptually different sources of inequality: the heterogeneity of cognitive augmentation itself, captured by $\beta(i,t)$, and the equilibrium amplification of these differences through production reallocation. Two economies characterized by similar distributions of AI-induced productivity gains may consequently display markedly different wage and employment responses simply because tasks differ in their degree of substitutability.

This distinction also helps reconcile apparently conflicting empirical findings. Some studies report substantial productivity gains with only modest changes in employment or wage inequality, whereas others document pronounced task reallocation following AI adoption. Within the present framework, these outcomes need not reflect differences in AI technologies themselves. They may instead arise because comparable improvements in cognitive augmentation propagate through production systems characterized by different values of σ. When tasks are highly substitutable, local productivity gains are rapidly transformed into reallocations of labor demand and sharper wage differentials. When complementarities prevail,

the same productivity improvements diffuse more broadly across production, generating smaller distributional adjustments despite similar technological progress.

The interaction between σ and the covariance structure discussed in section 4 further reinforces this point. The aggregate consequences of AI depend not only on the distribution of cognitive augmentation across tasks, but also on whether these productivity gains accrue primarily to tasks that are already highly productive. Estimating the covariance between baseline productivity and cognitive augmentation is therefore informative only once the elasticity of substitution governing the propagation of these gains has been identified.

The empirical counterpart of the model is thus naturally sequential. Once the local effects of AI on effective labor productivity have been characterized through $\beta(i,t)$ and $\lambda(i)$, estimating σ allows these microeconomic responses to be mapped into aggregate outcomes. The remaining question concerns whether higher effective labor productivity ultimately delays automation or is eventually dominated by improvements in AI capability. Addressing this issue requires identifying the evolution of automation capability relative to augmented labor productivity, the next structural parameter of the framework.

Recalling the definition of task-specific productive capability of AI capital per unit of financing cost, $\gamma(i,t) = \frac{\alpha_K(i,t)}{R(t)}$, automation occurs whenever the productivity per unit cost of AI-powered capital exceeds that of effective labor. At the automation frontier,

$$\gamma(i,t) = \frac{\alpha_L(i)\beta(i,t)}{W(i,t)}. \tag{38}$$

Equation (38) highlights an important implication of the model. Improvements in AI capability do not translate mechanically into automation. Their effect depends on the simultaneous evolution of cognitive augmentation and equilibrium wages. Faster improvements in AI capability increase the incentive to automate, whereas stronger growth in effective labor productivity raises the relative competitiveness of labor, delaying the expansion of the automation frontier.

Taking logarithms and first differences yields:

$$\Delta \log \gamma(i,t) + \Delta \log W(i,t) - \Delta \log \beta(i,t) \lesseqgtr 0, \tag{39}$$

which summarizes the conditions under which automation becomes more or less likely. The relevant empirical object is therefore automation pressure, namely the evolution of AI capability relative to augmented labor productivity, rather than AI exposure itself.

This distinction has important empirical implications. Existing studies typically proxy automation using measures of occupational AI exposure, implicitly assuming that technological capability is sufficient to explain labor displacement. Equation (39) suggests a different interpretation. Exposure identifies where AI may become economically relevant, whereas realized automation depends on the joint evolution of automation capability, cognitive augmentation, and equilibrium wages. Tasks experiencing rapid gains in effective labor productivity should therefore display systematically lower automation rates than equally exposed tasks where augmentation remains limited.

The model consequently suggests estimating automation through a structural hazard rather than a reduced-form exposure equation. Defining:

$$A(i,t) = \Delta \log \gamma(i,t) + \Delta \log W(i,t) - \Delta \log \beta(i,t), \tag{40}$$

the probability that task i becomes automated between periods $t-1$ and t can be written as:

$$Pr\big[Automation_{i,t} = 1\big] = \Lambda\big(A(i,t)\big), \tag{41}$$

where $\Lambda(\cdot)$ captures adoption frictions, organizational adjustment costs, and unobserved heterogeneity across firms and tasks. Unlike conventional hazard specifications, equation (41) identifies the structural determinants of automation implied by the model rather than relying on technological exposure alone.

This formulation yields a straightforward empirical prediction. Conditional on AI capability, automation hazards should be systematically lower where cognitive augmentation grows more rapidly or where productivity gains are reflected in higher equilibrium wages. Conversely, automation should be concentrated in tasks where improvements in AI capability consistently outpace the evolution of augmented labor productivity. Identifying these relative dynamics provides the empirical counterpart of the automation frontier developed in the theoretical framework.

Once the intensive margin of adjustment has been characterized through cognitive augmentation and automation, the remaining empirical question concerns the extensive margin: the extent to which AI also expands the frontier of economically valuable human activities. This is summarized by the task-creation parameter ψ, to which we now turn. In fact, unlike the standard automation-reinstatement framework, where new tasks are introduced to offset labor displacement, the present model links task creation directly to cognitive augmentation. As AI reduces the cost of prediction and information processing, activities that were previously uneconomical become productive because human judgment can be deployed more effectively. The expansion of the task frontier therefore arises from the same mechanism that raises effective labor productivity within existing tasks.

Let N_t denote the upper bound of the task frontier. Its evolution is given by:

$$\Delta N_t = \psi \sum_{i\in(I,N]} \lambda(i,t)\Delta\beta(i,t) + \nu_{t+1}, \tag{42}$$

where the summation extends over contemporaneous labor-performed tasks, $\psi > 0$ measures the responsiveness of task creation to cognitive augmentation, and ν_{t+1} captures technological and demand shocks unrelated to AI. The parameter ψ therefore summarizes the economy's capacity to transform improvements in prediction into new labor-intensive activities, providing the empirical counterpart of the reinstatement margin.

Equation (42) also highlights a distinctive implication of the AR–CA framework. Since cognitive augmentation depends on accumulated expertise through equation (16), task creation is itself shaped by the economy's ability to reproduce expertise over time. The expansion of the task frontier is therefore not determined exclusively by technological progress, but also by the organizational conditions governing learning and human capital accumulation. The apprenticeship mechanism consequently influences both the productivity of existing tasks and the emergence of new ones, linking short-run labor-market adjustment to long-run structural transformation.

The empirical implications follow directly. Estimating ψ requires observing the emergence of new economically relevant activities rather than simply changes in employment within existing occupations. Administrative occupation codes are therefore unlikely to be sufficient, since many AI-enabled tasks initially appear within established occupations before evolving into distinct organizational roles. Vacancy data, detailed task descriptions, online job postings, and within-occupation measures of task composition provide a more suitable empirical basis for identifying the expansion of the task frontier and distinguishing technological environments in which AI

primarily reallocates existing work from those in which it persistently generates new productive activities.

The productivity of these newly created tasks ultimately depends on the stock of expertise available to exploit cognitive augmentation. Assessing the long-run consequences of AI therefore requires one final parameter describing the efficiency with which expertise itself is accumulated χ, which captures the extent to which firms internalize the future returns generated by apprenticeship.

As discussed in section 3.2, expertise constitutes the complementary input through which workers transform AI-generated prediction into productive judgment. Unlike physical or digital capital, however, expertise is accumulated through production itself. Entry-level tasks jointly generate current output and future productive capacity by allowing workers to acquire tacit knowledge and contextual understanding subsequently embodied in cognitive augmentation. The evolution of expertise therefore depends on firms' hiring and training decisions, while its contribution to production is reflected in the augmentation factor $\beta(i,t)$. As a result, high values of χ characterize environments in which organizations capture most of the productivity gains associated with training. Lower values arise whenever worker mobility, knowledge spillovers, or incomplete contracting prevent firms from fully appropriating the expertise produced through apprenticeship. In this case, automation remains privately optimal while reducing investment in the accumulation of expertise below its socially efficient level. The resulting divergence between private and social incentives is precisely the apprenticeship externality identified above.

The empirical counterpart of this mechanism differs fundamentally from the previous parameters. Whereas β, σ, γ and ψ can be inferred from observed productivity, employment, or task dynamics, identifying χ requires measuring the institutions through which expertise is produced and retained. Apprenticeship intensity, hiring of inexperienced workers, internal promotion, job tenure, worker mobility, mentorship, and the persistence of learning-intensive occupations all become natural empirical counterparts of the appropriability mechanism embodied in the model. Cross-country differences in labor-market institutions or firm-level differences in organizational practices provide corresponding sources of identification.

The importance of χ extends beyond labor-market dynamics. Through its effect on expertise accumulation, it influences the future evolution of cognitive augmentation and, through equation (42), the economy's capacity to generate new judgment-intensive activities. The efficiency of AI adoption therefore depends not only on how rapidly firms automate existing production, but also on whether the organizational conditions required to reproduce expertise are preserved as automation progresses. The macroeconomic consequences of AI are consequently shaped as much by the institutions governing human capital formation as by the pace of technological innovation itself.

Taken together, the empirical agenda outlined above shifts attention from the incidence of AI adoption to the mechanisms through which it propagates across the economy. Exposure measures remain informative about the diffusion of technology and the activities most likely to be affected, while the structural parameters identified by the model characterize how those effects materialize. Cognitive augmentation governs productivity within human-performed tasks; the elasticity of substitution determines how local productivity gains propagate across production; automation capability shapes the boundary between human and machine activity; task creation governs the expansion of new productive opportunities; and the appropriability of expertise determines the efficiency of human capital accumulation over time. Since the underlying decisions on AI deployment, organizational design, and apprenticeship are taken within firms, their macroeconomic significance ultimately rests on how these choices interact through labor

reallocation, product markets, and expertise accumulation. Identifying these transmission channels provides an empirical basis for assessing their consequences for productivity, employment, wage inequality, factor shares, and long-run growth, as well as for determining which features of AI can be accommodated within established theories of technological change and which require new analytical margins.

Pursuing this agenda places correspondingly demanding requirements on data and identification. Estimation will require linking matched employer–employee records on employment histories, wages, promotions, training, and worker mobility to firm-level information on AI adoption, investment, and organizational practices; task-level evidence from vacancy postings, occupational surveys, workflow data, and enterprise software; and digital traces of AI use generated through coding platforms and interactions with large language models. Recent contributions by Aghion et al. (2025) and Lindenlaub et al. (2026) illustrate progress in this direction. Such data can support complementary econometric approaches – panel methods, quasi-experimental designs, and structural estimation – suited to recovering different components of the transmission process and ultimately mapping firm-level AI adoption into its aggregate economic effects.

7. Conclusion

The starting point of this paper is a simple observation. The emerging empirical evidence on generative AI does not fit comfortably within the traditional paradigms used to study technological change and labor markets. Early adjustments are concentrated among entry-level workers rather than along conventional skill dimensions; AI simultaneously substitutes for some cognitive activities while complementing others; and organizational change appears to play a role at least as important as technological capability itself. These patterns suggest that understanding AI requires moving beyond the view of technological progress as either a uniform increase in factor productivity or a process of routine automation alone.

The framework developed here provides one possible way of organizing this emerging evidence. By integrating the automation-reinstatement approach with the decision-theoretic perspective on AI as a form of cognitive augmentation, it interprets AI as a technology operating simultaneously on two margins. It expands the range of tasks that can be performed autonomously by capital while increasing the productivity of human labor in activities where prediction continues to require judgment, interpretation, and contextual reasoning. The resulting interaction reshapes comparative advantage within firms, modifies the organization of production, and alters the process through which expertise itself is accumulated. In this environment, labor-market outcomes depend not on automation alone, but on the joint evolution of cognitive augmentation, task substitution, task creation, and expertise formation.

Viewed from this perspective, several empirical regularities documented in the recent literature acquire a common interpretation. The concentration of employment losses among entry-level workers, the distinction between automating and augmenting uses of AI, the limited short-run response of wages, and the coexistence of rapid technological diffusion with modest aggregate employment effects all emerge as different manifestations of the same underlying adjustment process. In particular, the framework rationalizes seniority-biased technological change as the equilibrium outcome of the interaction between task substitution and cognitive augmentation, while the apprenticeship externality highlights a dynamic channel through which

organizational decisions today influence the future accumulation of expertise and, consequently, the economy's capacity to benefit from AI tomorrow.

The analysis of balanced growth further clarifies the macroeconomic implications of this interaction. Stable labor shares need not rely on a slowdown in automation, nor solely on the endogenous creation of new tasks. Instead, balanced growth requires a sufficiently rapid expansion of labor-augmenting cognitive capabilities alongside automation and task creation, allowing AI to increase productivity without progressively eroding labor's role in production. Unlike previous general-purpose technologies, generative AI simultaneously operates on both sides of this adjustment, affecting the productivity of capital and the effectiveness of human cognition. Whether these two forces evolve in a sufficiently balanced manner is therefore likely to become one of the central macroeconomic questions raised by AI diffusion.

The broader contribution of the paper lies less in the specific model than in the research agenda it suggests. The AR-CA synthesis identifies a small set of structural parameters that govern the aggregate consequences of AI adoption: the intensity of cognitive augmentation, the elasticity of substitution across tasks, the evolution of automation capability, the responsiveness of task creation, and the appropriability of expertise accumulated through apprenticeship. Together, these parameters determine how local organizational changes propagate into aggregate productivity, employment, wage inequality, labor's share, and long-run growth. Future empirical work aimed at identifying these objects, rather than relying exclusively on measures of AI exposure, may therefore offer a more informative basis for understanding the macroeconomics of AI.

Several directions for future research follow naturally. A first extension would endogenize the direction of AI-driven technical change, allowing automation, cognitive augmentation, and task creation to evolve jointly in response to economic incentives and public policy. A second concerns the interaction between AI, market structure, and the ownership of complementary intangible assets, factors that are likely to influence both the pace of organizational adaptation and the distribution of rents generated by AI. Finally, the framework points to the need for richer empirical evidence capable of jointly measuring task reallocation, cognitive augmentation, expertise accumulation, and apprenticeship dynamics within firms. Such evidence would make it possible not only to quantify the mechanisms emphasized here, but also to assess a broader question that remains largely unanswered: whether the long-run efficiency of AI adoption ultimately depends as much on the economy's ability to produce expertise as on its ability to automate production.

References

Abraham K. and Kerney S. (2020), "Explaining the decline in the US employment-to-population ratio: a review of the evidence", *Journal of Economic Literature*, 58, pp. 585-643.

Acemoglu D. (2025), "The simple macroeconomics of AI", *Economic Policy*, 40, pp. 13-58.

Acemoglu D. and Autor D. (2011), "Skill, tasks and technology: implications for employment and earnings", in Ashenfelter O. and Card D. (eds), *Handbook of Labor Economics*, *Vol. 4, Part A*, Amsterdam: Elsevier.

Acemoglu D., Autor D., Hazell J. and Restrepo P. (2022), "Artificial intelligence and jobs: evidence from online vacancies", *Journal of Labor Economics*, 40, pp. S293-S340.

Acemoglu D., Autor D. and Johnson S. (2026), "Building pro-worker artificial intelligence", *NBER Working Paper*, no. 34854, Cambridge (MA): National Bureau of Economic Research.

Acemoglu D. and Restrepo P. (2018), "The race between man and machine; implications of technology for growth, factor shares, and employment", *American Economic Review*, 108, pp. 1488-1542.

Acemoglu D. and Restrepo P. (2019a), "Automation and new tasks: how technology displaces and reinstates labor", *Journal of Economic Perspectives*, 33, pp. 3-30.
Acemoglu D. and Restrepo P. (2019b), "Artificial intelligence, automation and work", in Agrawal A., Gans J. and Goldfarb A. (eds), *The Economics of Artificial Intelligence: An Agenda* (pp. 197-236), Chicago: University of Chicago Press.
Acemoglu D. and Restrepo P. (2022), "Task, automation and the rise in U.S. wage inequality", *Econometrica*, 90, pp.1973-2016.
Aghion P., Jones B. and Jones C. (2019a), "Artificial intelligence and economic growth", in Agrawal A., Gans J. and Goldfarb A. (eds), *The Economics of Artificial Intelligence: An Agenda* (pp. 237-290), Chicago: University of Chicago Press.
Aghion P., Bergeaud A., Blundell R. and Griffith R. (2019b), "The innovation premium to soft skills in low-skilled occupations", *Banque de France Working Paper*, no. 739, Paris: Banque de France. Available online.
Aghion P., Bunel S., Jaravel X., Mikaelsen T., Roulet A. and Søgaard J. (2025), "How different uses of AI shape labor demand: evidence from France", *American Economic Review Papers and Proceedings*, 115, pp. 62-67.
Agrawal A., Gans J. and Goldfarb A. (2018), *Prediction Machines: The Simple Economics of Artificial Intelligence*, Cambridge (MA): Harvard Business Press.
Agrawal A., Gans J. and Goldfarb A. (2019a), "Artificial Intelligence: the ambiguous labor market impact of automating prediction", *Journal of Economic Perspectives*, 33, pp. 31-50.
Agrawal A., Gans J. and Goldfarb A. (2019b), "Prediction, judgment and complexity: a theory of decision-making and artificial intelligence", in Agrawal A., Gans J. and Goldfarb A. (eds), *The Economics of Artificial Intelligence: An Agenda* (pp. 89-110), Chicago: University of Chicago Press.
Agrawal A., Gans J. and Goldfarb A. (2025), "The economics of bicycles for the mind", *NBER Working Paper*, nr. 34034, Cambridge (MA): National Bureau of Economic Research.
Agrawal A., McHale J. and Oettl A. (2026), "Enhancing worker productivity without automating tasks: a different approach to AI and the task-based model", *NBER Working Paper*, no. 34781, Cambridge (MA): National Bureau of Economic Research.
Althoff L. and Reichardt H. (2026), "Task-specific technical change and comparative advantage", mimeo.
Akin M. (2026), "The apprenticeship externality: generative AI, entry-level work, and the future supply of expertise", SSRN. Available online.
Autor D. (2013), "The 'task approach' to labor markets: an overview", *Journal of Labor Market Research*, 46, pp. 185-199.
Autor D., Dorn D., Katz L., Patterson C. and Van Reenen J. (2020), "The fall of the labor share and the rise of the superstar firms, *Quarterly Journal of Economics*, 135, pp. 645-709.
Autor D., Levy F. and Murnane R. (2003), "The skill content of recent technological change: an empirical exploration", *Quarterly Journal of Economics*, 118, pp. 1279-1333.
Autor D. and Salomons A. (2018), "Is automation labor share–displacing? Productivity growth, employment, and the labor share", *Brookings Papers on Economic Activity*, 1-63.
Autor D. and Thompson N. (2025), "Expertise", *Journal of the European Economic Association*, 23, pp. 1203-1271.
Bresnahan T. and Trajtenberg M. (1995), "General purpose technologies: engines of growth?", *Journal of Econometrics*, 65, pp. 83-108.
Brynjolfsson E., Chandar K. and Chen R. (2025a), "Canaries in the coal mine: six facts about the recent employment effects of artificial intelligence", mimeo, Stanford Digital Economy Lab. Available online.
Brynjolfsson E., Li D. and Raymond L. (2025b), "Generative AI at work", *Quarterly Journal of Economics*, 140, pp. 889-942.
Brynjolfsson E. and McAfee A. (2014), *The Second Machine Age: Work, Progress and Prosperity in a Time of Brilliant Technologies*, New York: W.W. Norton & Company.
Cazzaniga M., Jaumotte F., Li L., Melina G., Panton A., Pizzinelli C., Rockall E. and Tavares M. (2024), "Gen-AI: artificial intelligence and the future of work", *IMF Staff Discussion Note*, no. 2024/001, Washington (DC): International Monetary Fund. Available online.
Chanoi C. and Banger-Drowns C. (2025), "AI exposure by U.S. occupations and work tasks and the effect on wages", *Washington Center for Equitable Growth Working Paper*, no. 102325, Washington (DC): Washington Center for Equitable Growth. Available online.
Crafts N. (2021), "Artificial Intelligence as a general-purpose technology: an historical perspective", *Oxford Review of Economic Policy*, 37, pp. 521-536.
Eloundou T., Manning S., Mishkin P. and Rock P. (2024), "GPTs are GPTs: labor market impact potential of LLMs", *Science*, 384, pp. 1306-1308.
Emmanuel N. Harrington E. and Pallais A. (2026), "The power of proximity to coworkers", *Quarterly Journal of Economics*, 141, pp. 1825-1870.

Felipe J., Bajaro D.F., Estrada G. and McCombie J. (2020), "What do tests of the relationship between employment growth and technical progress hide?", *PSL Quarterly Review*, 73, pp. 367-392.
Filippucci F., Gal P., Laengle K., Schief M. and Yildirim M. (2026), "AI meets trade: global linkages and the cross-country distribution of the gains from AI", *OECD Artificial Intelligence Papers*, no. 57, Paris: OECD Publishing. Available online.
Frey C. and Osborne M. (2017), "The future of employment: how susceptible are jobs to computerization?", *Technological Forecasting and Social Change*, 114, pp. 254-280.
Garicano L. and Rayo L. (2026), "Training in the age of AI: a theory of apprenticeship viability", *CEPR Discussion Paper*, no. DP20634, Paris and London: CEPR Press.
Goldin C. and Katz L. (2008), *The Race between Education and Technology*, Cambridge (MA): Harvard University Press.
Goller D., Gschwendt C. and Wolter S. (2025), "This time it's different. Generative artificial intelligence and occupational choice", *Labour Economics*, 95, art. 102746.
Goos M., Manning A. and Salomons A. (2014), "Explaining job polarization: routine-biased technological change and offshoring", *American Economic Review*, 104, pp. 2509-2526.
Greenan N., Guarascio D. and Reljic J. (2025), "AI and the labour market: opening the black box", *Eurasian Business Review*, 15, pp. 925-951.
Grossman G, and Oberfield E. (2022), "The elusive explanation for the declining labor share", *Annual Review of Economics*, 32, pp. 93-124.
Guarascio D. and Reljic J. (2025), "AI and employment in Europe", *Economics Letters*, 247, art. 112183.
Hémous D. and Olsen M. (2022), "The rise of the machines: automation, horizontal innovation, and income inequality", *American Economic Journal: Macroeconomics*, 14, pp. 179-223.
Hosseini S. and Lichtinger G. (2025), "Generative AI as seniority-biased technological change: evidence from U.S. résumé and job posting data", SSRN. Available online.
Humlum A. and Vestergaard E. (2025), "Large language models, small labor markets effects", *NBER Working Paper*, no. 33777, Cambridge (MA): National Bureau of Economic Research. Available online.
Ide E. (2026), "Automation, AI, and the intergenerational transmission of knowledge", arXiv:2507.16078. Available online.
Iscenko Z. and Millet F. (2026), "Looking for the ladder: is AI impacting entry-level jobs?", Economic Innovation Group Commentary, The American Worker Project Working Paper. Available online.
Jakubik A., Rotunno L. and Saini A. (2025), "Foresee the unseen: evaluating the impact of artificial intelligence on international trade", *Economic Modeling*, 47, pp. 842-861.
Jones C. (2024), "The A.I. dilemma: growth versus existential risk", *American Economic Review: Insights*, 6, pp. 575-590.
Jones C. and Tonetti C. (2026), "Past automation and future AI: how weak links tame the growth explosion", mimeo. Available online.
Kaldor N. (1961), "Capital accumulation and economic growth", in Hague D.C. (ed.), *The Theory of Capital* (pp. 177-222), London: Palgrave Macmillan. Available online.
Katz L. and Murphy K. (1992), "Changes in relative wages, 1963-1987: supply and demand factors", *Quarterly Journal of Economics*, 107, pp. 35-87.
Korinek A. and Juelfs M. (2022), "Preparing for the (non-existent?) future of work", in Bullock J., Chen Y.-C., Himmelreich J., Hudson V., Korinek A., Young M. and Zhang B. (eds), *The Oxford Handbook of AI Governance* (pp. 746-776), London and New York: Oxford University Press.
Lambert P. and Schindler Y. (2026), "The broken ladder: AI, remote work, and early-career hiring", *CAGE Online Working Paper Series*, no. 808/2026, Warwick: University of Warwick, Competitive Advantage in the Global Economy. Available online.
Lindenlaub I., Oh R., Rodriguez M. and Veldkamp L. (2026), "Beyond exposure: predicting AI adoption based on comparative advantage", *CEPR Discussion Paper*, no. 21589, Paris and London: CEPR Press.
Moffitt R. (2012), "The reversal of the employment-population ratio in the 2000s: facts and explanations", *Brookings Papers on Economic Activity*, 2, pp. 201-264.
Morandini S., Fraboni F., De Angelis M., Puzzo G., Giusino D. and Pietrantoni L. (2023), "The impact of artificial intelligence on workers' skills: Upskilling and reskilling in organisations", *Informing Science*, 26, pp. 39-68.
Nigar M., Juli J., Golder U., Alam M. and Hossain M. (2025), "Artificial intelligence and technological unemployment: understanding trends, technology's adverse roles, and current mitigation guidelines", *Journal of Open Innovation: Technology, Market, and Complexity*, 11, art. 100607.
OECD – Organisation for Economic Co-operation and Development (2024), *Digital Economy Outlook 2024*, Paris: OECD Publishing.
Pizzinelli C., Panton A., Mendes Tavares M., Cazzaniga M. and Li L. (2023), "Labor market exposure to AI: cross-country differences and distributional implications", *IMF Working Paper*, nr. 2023/216, Washington (DC): International Monetary Fund. Available online.

Poláková M., Suleimanová J., Madzík P., Copuš L., Molnárova I. and Polednová J. (2023), "Soft skills and their importance in the labour market under the conditions of Industry 5.0", *Heliyon*, 9, art. e18670.
Schubert G. (2025), "Organizational technology ladders: remote work and generative AI adoption", SSRN. Available online.
The Financial Times (2025), "AI will lead one in four big UK businesses to cut staffing, research shows", November 9.
Cain Miller C. and Cox C. (2023), "In reversal because of A.I., office jobs are now more at risk", *The New York Times*, August 30. Available online.
Trammell P. and Korinek A. (2025), "Economic growth under transformative AI", *NBER Woking Paper*, no. 31815, Washington (DC): National Bureau of Economic Research.
Zarifhonarvar A. (2024), "Economics of ChatGPT: a labor market view on the occupational impact of artificial intelligence", *Journal of Electronic Business and Digital Economics*, 3, pp. 100-116.
Zeira J. (1998), "Workers, machines, and economic growth", *Quarterly Journal of Economics*, 113, pp. 1091-1117.