The Labor Market Effects of AI Automation and Augmentation*

Tomáš Oleš† Francesco Roncone†

September 21, 2026

Abstract

Does artificial intelligence automate work or augment workers' productive capabilities? Our results show that it can do both, through distinct channels that conventional exposure measures obscure. We distinguish these channels by linking U.S. patents to occupations' tasks and tools, with patent–task matches capturing automation and patent–tool matches capturing augmentation. Using the origins and diffusion of historical AI breakthroughs to predict contemporary local exposure, we find large but opposing labor-market effects. A one-standard-deviation increase in AI automation reduces occupational employment by 21 percent, with little effect on wages; AI augmentation instead raises wages by 8 percent, alongside some employment growth. We also find that a larger AI share within each channel strengthens labor displacement through automation while shifting augmentation's effects from employment growth toward wage gains. The labor-market consequences of AI therefore depend critically on the balance between automation and augmentation rather than overall exposure alone.

Keywords: Artificial Intelligence; Automation; Augmentation; Local Labor Markets; Occupational Tasks; Occupational Tools.

JEL Codes: J23, J24, J30, O33.

*We thank Martin Lábaj, Klaus Prettner, Magnus Lodefalk, Ekaterina Prytkova, Giulio Zanella, Qazi Haque, and Tomáš Ševček for helpful comments and discussions that improved the paper. We also thank participants at the SGH Warsaw School of Economics Workshop, the Manic Mondays seminar series at Sapienza Università di Roma, the SEA–NOeG Joint Conference in Vienna, and the Workshop on Technology, Labor Markets, and Shared Prosperity at Bratislava University of Economics and Business for valuable feedback. We acknowledge financial support from research project APVV-23-0090, "Towards a Shared Prosperity: Technology, Inequality, and Labor Market," and from the European Union through Slovakia's Recovery and Resilience Plan under project 09I03-03-V04-00516, "Fellowships for Excellent Researchers R2–R4." All remaining errors are our own.

†Department of Public Policy, Bratislava University of Economics and Business, Dolnozemská cesta 1/b, 852 35 Bratislava, Slovakia.

1 Introduction

The accelerating momentum of artificial intelligence (AI) has intensified concerns about its impact on the labor market. A central question is whether AI will primarily replace human labor by performing tasks autonomously or enhance workers' capabilities and boost their productivity. Consider two workers in different jobs exposed to the same recent wave of AI: a *statistical assistant* increasingly relies on AI software that automatically classifies, organizes, and retrieves information that previously required manual processing. A *radiologist*, by contrast, uses AI-assisted diagnostic software to interpret medical images more quickly and accurately while retaining responsibility for the diagnosis. In both cases AI changes how work is performed, yet in fundamentally different ways. In the first, AI substitutes for human labor by performing workers' tasks; in the second, it complements labor by improving the tools through which work is carried out.

Understanding the distinct roles of these two forces is therefore crucial for predicting AI's impact on employment and wages and, ultimately, whether AI helps or harms workers. Although fears of large-scale labor displacement have concerned economists and workers for decades (see Brynjolfsson and McAfee 2015), historical experience and modern task-based theory suggest that automation has often been counterbalanced by the creation of new labor-intensive tasks, leading to a reshuffling of workers across tasks rather than mass unemployment (Acemoglu and Restrepo 2018; Acemoglu and Restrepo 2019). AI, however, may differ from earlier waves of technological change because it can potentially automate a much broader range of activities (Acemoglu 2021; Acemoglu et al. 2026), including non-routine cognitive tasks in which human labor has increasingly specialized in response to the automation of routine work over recent decades (e.g., Autor et al. 2003; Autor and Dorn 2013). At the same time, AI can operate as a worker-augmenting tool, raising productivity in tasks that remain assigned to labor (Agrawal et al. 2026). Which of these channels predominates in AI-driven technological change—and with what consequences for employment and wages—remains an open empirical question.

Answering this question requires distinguishing between the two channels. A single measure of AI exposure may mask their potentially opposing effects and obscure their distinct consequences for workers. We operationalize these channels empirically through a simple task–tool distinction: tasks describe what workers do to produce output, whereas tools describe what workers use to perform their tasks. Innovations that resemble an occupation's tasks are therefore more likely to automate activities currently performed by workers, whereas innovations that resemble its tools are more likely to augment workers by enhancing their productive capabilities. We implement this approach using transformer-based language models that measure the semantic similarity

between the text of each U.S. patent and the task and tool descriptions of detailed O*NET occupations. Patent–task similarity defines our measure of automation exposure, whereas patent–tool similarity defines our measure of augmentation exposure. Combining these measures with the USPTO's Artificial Intelligence Patent Dataset (Pairolero et al. 2025) separates each channel into its AI and non-AI components. This allows us to examine both how AI-driven automation and augmentation affect labor-market outcomes and whether the AI share within each channel matters beyond workers' overall exposure to automating and augmenting innovation.

Although our patent–task and patent–tool similarity measures are derived from the entire universe of U.S. utility patents, a second key feature of our approach is that we explicitly incorporate a local dimension into the measurement of technological exposure. Building on evidence that knowledge spillovers are geographically localized (Jaffe et al. 1993) and that adjustments in employment and skills associated with new technologies spread only gradually across space (Kalyani et al. 2025), we posit that workers are more directly exposed to innovations located in their surrounding economic environment. Consequently, technological change may affect even the same occupation differently across local labor markets. For example, a chemical technician employed in a commuting zone with intensive laboratory-instrument innovation is likely to experience substantially greater exposure to productivity-enhancing augmentation than an otherwise similar technician elsewhere. We capture this geographic variation by assigning patents to commuting zones based primarily on assignee locations and constructing occupation–commuting-zone exposure measures. Although we do not observe workplace adoption directly, patent assignees—most of which are firms—provide a natural geographic anchor for these innovations.

A central challenge for identifying the causal effects of local AI exposure on labor-market outcomes is that innovation is not randomly distributed across space. To address this concern, we instrument contemporaneous local AI automation and augmentation exposure using measures designed to capture the spatial diffusion of historical AI breakthroughs originating in more distant commuting zones. Specifically, we use the classification developed by Kelly et al. (2021) to identify patents granted between 1980 and 2000—well before the recent wave of widespread AI adoption—that are textually distinct from prior inventions but closely related to subsequent inventions, thereby capturing novel technological directions on which later innovation builds. We combine this classification with the USPTO's Artificial Intelligence Patent Dataset to identify AI-related breakthroughs, determine their automation and augmentation relevance based on our patent–task and patent–tool similarity scores, and weight them by geographic distance. To reduce concerns that the instruments reflect persistent local conditions or direct spillovers from nearby innovation,

we exclude breakthroughs originating in or near the focal commuting zone. We further show that our estimates are robust to a range of alternative instrument constructions and are not driven by breakthroughs originating in a small number of historically dominant AI innovation hubs, such as Silicon Valley or the Northeast Corridor.

The analysis yields three main findings. First, distinguishing between AI-driven automation and augmentation reveals markedly different labor-market effects. Our preferred 2SLS estimates imply that, between 2005–07 and 2022–24, a one-standard-deviation increase in AI automation exposure reduces occupational wage bills by approximately 22 percent and employment by 21 percent, with little effect on hourly wages. By contrast, an equivalent increase in AI augmentation exposure raises occupational wage bills by approximately 16 percent, reflecting hourly-wage gains alongside employment growth. This contrast between the effects of automation and augmentation is particularly pronounced in recent years.

Second, the composition of technological change matters in addition to its overall intensity. Holding total automation and augmentation exposure constant, we find that a greater AI share within automating innovation strengthens employment displacement, whereas a greater AI share within augmenting innovation shifts adjustment away from employment growth and toward wage gains. Together, these findings are consistent with concerns that AI may reinforce the labor-displacing direction of technological change (Acemoglu and Restrepo 2020; Acemoglu 2021), while also illustrating its potential to expand workers' productive capabilities when it augments rather than automates human labor (Acemoglu 2025; Acemoglu et al. 2026).

Third, the effects of AI depend importantly on the workers and tasks exposed to it. More highly educated local workforces experience larger employment losses when innovation shifts toward AI automation, but benefit more when that shift occurs through augmentation. Routine-intensive occupations, by contrast, fare worse as innovation becomes more AI-intensive, particularly through the automation channel. These patterns highlight that who gains and who loses from AI depends not only on exposure itself, but also on workers' skills, the tasks they perform, and the kind of AI they face.

Our paper contributes to the growing empirical literature on the labor-market consequences of artificial intelligence. Existing approaches commonly summarize occupational AI exposure using a single index based on relationships between AI and occupational tasks or abilities (e.g., Brynjolfsson et al. 2018; Webb 2019; Felten et al. 2021). Combining these measures with online job vacancy data, Acemoglu et al. (2022) document reductions in non-AI job postings and changes in skill requirements at AI-exposed establishments, but find no detectable employment or wage

effects in more AI-exposed occupations and industries. More recent empirical evidence paints a more nuanced picture, documenting substantial productivity gains in particular workplace settings (Brynjolfsson et al. 2025; Cui et al. 2026), changes in task content within occupations and the reallocation of labor across tasks within firms (Gathmann et al. 2024; Hampole et al. 2025), limited short-run effects on earnings and hours despite widespread adoption and task reorganization (Humlum and Vestergaard 2025), and relative employment declines among younger workers in highly AI-exposed occupations (Brynjolfsson et al. 2026). Together, this evidence shows that AI is already reshaping work. However, because existing studies typically summarize exposure along a single dimension or focus on particular adoption settings, the relative roles of automation and augmentation remain empirically unsettled. We address this gap by bringing both channels into a unified task–tool measurement framework and estimating their distinct labor-market effects across occupations and local labor markets throughout the United States.

Our study is most closely related to an emerging literature that distinguishes between labor-automating and labor-augmenting technological change. Autor et al. (2024) identify automating innovations through their relationship with occupational tasks and augmenting innovations through their relationship with job-title-based descriptions of occupational outputs, while Marguerit (2025) extends this approach specifically to AI. Kogan et al. (2023) and Chen and Wang (2024) instead distinguish labor-substituting from labor-augmenting technologies using different classification strategies. The former matches patents to occupational tasks classified as routine or non-routine, interpreting the resulting exposures as labor-saving or labor-augmenting, respectively. The latter groups AI patents by functional capability and then labels the resulting categories as augmenting or displacing based on their estimated employment effects within exposed occupational segments. Compared with these approaches, our study differs along three dimensions. First and most importantly, we identify augmentation through innovations related to the tools workers use. This complements Autor et al. (2024) and Marguerit (2025) by capturing a distinct margin of augmentation that may be particularly relevant for AI because improved tools can enhance workers' productive capabilities within established occupations without necessarily changing occupational outputs. Moreover, our approach does not impose an exclusive ex ante classification, and the two channels are defined independently of subsequent labor-market outcomes. Consequently, the same innovation may naturally affect a given occupation through automation, augmentation, or both. Second, we distinguish AI from non-AI innovation within both channels, allowing us to test whether the AI content of technological change matters beyond the overall intensity of automating and augmenting innovation. To our knowledge, ours is

the first study to examine this AI-versus-non-AI composition within both channels. Third, we introduce a spatial dimension by allowing the same occupation to face different levels of technological exposure across local labor markets, consistent with evidence that technological diffusion is geographically localized and gradual (Jaffe et al. 1993; Kalyani et al. 2025) and extending recent work on the labor-market effects of AI and emerging digital technologies across U.S. regions (Huang 2024; Prytkova and Petit 2026).

The remainder of the paper proceeds as follows. Section 2 develops the conceptual framework. Section 3 describes the construction of the automation and augmentation exposure measures, and Section 4 documents their descriptive patterns. Section 5 presents the identification strategy, while Section 6 reports the main results and heterogeneity analyses. Section 7 concludes.

2 Conceptual Framework

We motivate the empirical analysis with a task-based framework that allows technological exposure to vary across local labor markets and distinguishes task automation, in which technologies perform tasks without labor, from tool augmentation, in which technologies enhance the tools used jointly with workers. Accordingly, the same occupation may face different combinations of automation and augmentation across local labor markets. The framework is deliberately technology-neutral: both AI and non-AI innovations may operate through either channel. Its purpose is to discipline this task–tool distinction and clarify its labor-market implications rather than to provide a separate structural theory of AI.[1]

2.1 Tasks and Tools in Production

Consider occupation o in commuting zone c. Occupational output is produced by combining a continuum of tasks indexed by $i \in [0, 1]$:

$$Y_{oc} = \left[\int_0^1 y_{oc}(i)^{\frac{\sigma-1}{\sigma}} \, di \right]^{\frac{\sigma}{\sigma-1}}, \qquad \sigma > 0, \quad \sigma \neq 1, \tag{1}$$

where σ denotes the elasticity of substitution across tasks. We leave the degree of substitutability across tasks otherwise unrestricted, as the conditional technological effects on which we focus

[1] The framework builds on the task-based approach of Acemoglu and Autor (2011) and Acemoglu and Restrepo (2018), in which production takes place through tasks that may be allocated between labor and capital. We draw in particular on the extension of Acemoglu and Restrepo (2018) in which labor-performed tasks combine labor with labor-complementary capital, while separate capital inputs can perform tasks autonomously. In the context of AI, Acemoglu (2025) similarly distinguishes extensive-margin automation from task complementarities, while Agrawal et al. (2026) emphasize the potential for AI to raise worker productivity without automating the tasks they perform.

below do not depend on a particular value of σ.

A task may be produced either by workers using complementary tools or, when technologically feasible, autonomously using capital, software, or algorithms. We represent output under the two production methods as

$$y_{oc}^{L}(i) = \gamma_{o}^{L}(i)\ell_{oc}(i)^{1-\kappa}\left\{\Gamma_{oc}x_{oc}(i)\right\}^{\kappa}, \qquad y_{oc}^{K}(i) = A_{oc}^{K}\gamma_{o}^{K}(i)k_{oc}(i), \qquad 0 < \kappa < 1. \tag{2}$$

Here, $\ell_{oc}(i)$ denotes labor and $x_{oc}(i)$ the occupational tool input used jointly with workers. The parameter Γ_{oc} captures the productivity of these labor-complementary tools, while κ governs the role of the tool input in the labor–tool production method. A common productivity shifter for the labor–tool method is normalized to one. By contrast, $k_{oc}(i)$ denotes autonomous capital, software, or algorithms, while A_{oc}^{K} captures their productivity. The functions $\gamma_{o}^{L}(i)$ and $\gamma_{o}^{K}(i)$ capture task-specific productivity under the two production methods.[2] Following the standard task framework, we order tasks such that labor's comparative advantage increases with the task index,

$$\frac{d}{di}\left[\frac{\gamma_{o}^{L}(i)}{\gamma_{o}^{K}(i)}\right] > 0. \tag{3}$$

For tasks for which both production methods are technologically feasible, the lower-cost production method is used. Let w_{oc} denote the occupational wage, R_{oc}^{X} the unit cost of the tool input, and R_{oc}^{K} the rental cost of autonomous capital. Cost minimization implies

$$q_{oc}^{L}(i) = \frac{1}{\gamma_{o}^{L}(i)}\left(\frac{w_{oc}}{1-\kappa}\right)^{1-\kappa}\left(\frac{R_{oc}^{X}}{\kappa\Gamma_{oc}}\right)^{\kappa}, \qquad q_{oc}^{K}(i) = \frac{R_{oc}^{K}}{A_{oc}^{K}\gamma_{o}^{K}(i)}. \tag{4}$$

For simplicity, we summarize technological feasibility and cost-minimizing assignment by an effective automation frontier $I_{oc} \in [0,1]$. Tasks below I_{oc} are performed autonomously, while tasks above I_{oc} are performed through the labor–tool method. A task may remain above the frontier either because autonomous production is not yet technologically feasible or because it is feasible but more costly than the labor–tool method. Appendix I separates these technological-feasibility and cost-minimization components of the effective frontier.

[2]The functions $\gamma_{o}^{L}(i)$ and $\gamma_{o}^{K}(i)$ are taken to be common across commuting zones, consistent with our use of occupational task and tool inventories from O*NET. Local variation instead enters through the technology shifters A_{oc}^{K} and Γ_{oc}, input costs, and the resulting task assignment.

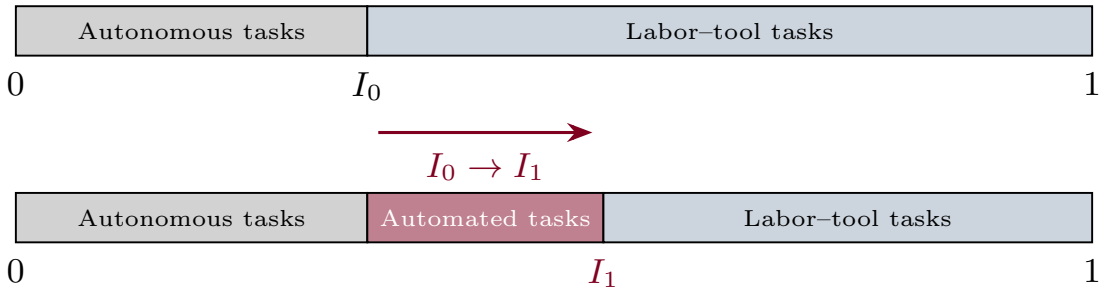
2.2 Automation and Augmentation

Within this environment, automation and augmentation capture two distinct channels through which technological change can affect work:

$$\underbrace{dI_{oc} > 0}_{\text{task automation}} \quad , \qquad \underbrace{d\Gamma_{oc} > 0 \;\text{ or }\; dR^X_{oc} < 0}_{\text{tool augmentation}} . \tag{5}$$

Automation expands the set of tasks performed autonomously. This may occur because technological progress makes autonomous production feasible for tasks that previously required workers or because it makes autonomous production cost-effective for tasks that were already technologically feasible. In either case, an increase in I_{oc} shifts additional tasks from the labor–tool method to autonomous production. Tool augmentation instead improves the complementary technology used by workers. Conditional on a task remaining performed through the labor–tool method, a higher Γ_{oc} raises the productivity of the tool input, while a lower R^X_{oc} reduces its unit cost. Both changes reduce the cost of producing the task using labor and tools. Figure 1 visualizes the two technological margins.

Panel A. *Task automation*



Panel B. *Tool augmentation*

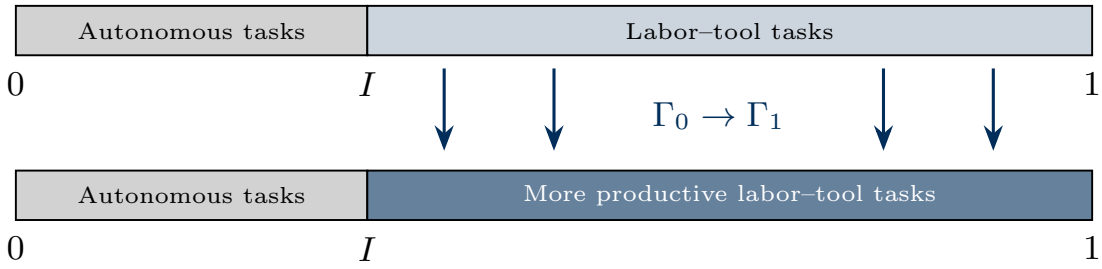


Figure 1: Task Automation and Tool Augmentation in Task Space

Notes: Panel A illustrates task automation as an increase in the effective automation frontier from I_0 to $I_1 > I_0$, shifting tasks between I_0 and I_1 from the labor–tool method to autonomous production. Panel B illustrates the Γ margin of tool augmentation as an increase from Γ_0 to $\Gamma_1 > \Gamma_0$, raising the productivity of tasks that continue to be performed using labor and tools while holding task allocation fixed.

The distinction shown in Figure 1 is particularly relevant for a powerful general-purpose technology like AI because its broad and increasingly flexible capabilities allow it to operate through different margins across tasks, workplace settings, and organizational structures. While the same AI technology may perform some activities autonomously, it can support workers in others by providing information, recommendations, software, or other complementary inputs that improve how workers perform their tasks.

Our formal definition of tool augmentation deliberately isolates this direct labor-complementary margin. More generally, improvements in occupational tools may also affect task assignment through relative costs or induce broader changes in the organization or composition of work, including the tasks workers perform. For simplicity, the comparative statics presented below hold task assignment fixed, while such adjustments may form part of the overall empirical labor-market response to tool-augmenting innovation.

2.3 From Technological Change to Labor-Market Outcomes

The two technological channels have distinct direct effects on labor. For a task produced using labor and tools, cost minimization implies

$$\ell_{oc}(i) = \frac{1-\kappa}{w_{oc}} q_{oc}^{L}(i) y_{oc}(i). \tag{6}$$

Aggregating over tasks performed by labor gives conditional occupational labor demand,

$$L_{oc} = \frac{1-\kappa}{w_{oc}} \int_{I_{oc}}^{1} q_{oc}^{L}(i) y_{oc}(i)\, di. \tag{7}$$

Automation. An outward shift in the automation frontier directly removes labor from the tasks that switch to autonomous production. Holding task output and factor prices fixed,

$$\left.\frac{\partial L_{oc}}{\partial I_{oc}}\right|_{\text{direct}} = -\frac{1-\kappa}{w_{oc}} q_{oc}^{L}(I_{oc}) y_{oc}(I_{oc}) < 0. \tag{8}$$

This is the direct displacement effect of automation. At the same time, lower production costs may expand occupational output, generating a countervailing scale effect. The equilibrium employment response therefore depends on the balance between the direct displacement effect, the countervailing scale effect, and associated factor-price adjustments.

Augmentation. Holding task assignment fixed, tool augmentation changes labor requirements within tasks that remain performed through the labor–tool method. Conditional on task output

and wages, equations (4) and (6) imply

$$\frac{\partial \ln \ell_{oc}(i)}{\partial \ln \Gamma_{oc}} = -\kappa, \qquad \frac{\partial \ln \ell_{oc}(i)}{\partial \ln R^{X}_{oc}} = \kappa, \qquad i > I_{oc}. \tag{9}$$

Hence, more productive tools or a lower cost of the tool input reduce the amount of labor required to produce a given quantity of a task. For given labor and tool inputs, a higher Γ_{oc} simultaneously raises the physical marginal product of labor. Letting $MPL^{L}_{oc}(i) \equiv \partial y^{L}_{oc}(i)/\partial \ell_{oc}(i)$,

$$\frac{\partial \ln MPL^{L}_{oc}(i)}{\partial \ln \Gamma_{oc}} = \kappa > 0, \qquad i > I_{oc}. \tag{10}$$

Tool augmentation can therefore raise worker productivity even as it reduces the labor input required per unit of task output. Whether employment rises or falls depends on the resulting equilibrium adjustment, including the extent to which lower production costs expand task and occupational output. Likewise, a higher physical marginal product of labor does not imply an unambiguously higher equilibrium wage, since wages depend on the value rather than the physical marginal product of labor and therefore also reflect adjustments in task prices and other equilibrium responses (Acemoglu 2025). At the occupation–commuting-zone level, these responses may also incorporate worker mobility and local price adjustments. Consequently, our framework delivers clear direct technological effects but leaves the equilibrium responses of employment, wages, and the wage bill to be determined empirically.

3 Measuring AI Automation and Augmentation

We translate the theoretical distinction between automation and augmentation into empirical measures of two channels of localized technological change, which we further separate into AI and non-AI components. The first, *automation exposure*, measures the extent to which local innovative activity is related to occupational tasks and therefore captures the potential for technology to substitute for labor. The second, *augmentation exposure*, measures the extent to which local innovative activity is related to the tools workers use in production and therefore captures the potential for technology to augment workers' productive capabilities.

3.1 AI Patent Classification

Our underlying data comprise the universe of U.S. utility patents granted between 1980 and 2021, including patent titles, abstracts, grant years, and locational information on inventors and assignees, obtained from the United States Patent and Trademark Office (USPTO) PatentsView

records. To distinguish AI-related from non-AI innovation, we rely on the Artificial Intelligence Patent Dataset (AIPD) 2023, developed by the USPTO. The dataset uses transformer-based language models trained on patent text to assign each patent a predicted probability for each of eight AI functions: machine learning, evolutionary computation, natural language processing, speech, computer vision, knowledge processing, AI hardware, and planning and decision systems (Pairolero et al. 2025). We classify a patent as AI-related if its predicted probability for at least one function is at least 86 percent, using the authors' midpoint threshold designed to balance precision and recall. Combining the AIPD classification with the patent–task and patent–tool similarity measures introduced below allows us to distinguish AI from non-AI automation and augmentation.

3.2 Patent–Task and Patent–Tool Similarity

Occupational task and tool inventories are obtained from O*NET version 25.0, which serves as the fixed reference for constructing our automation and augmentation measures. Appendix A.4 describes their assignment to our occupation panel, while Appendix Table A.1 provides examples from the resulting task and tool inventories for four occupations. We measure the textual similarity between patents and occupational task and tool descriptions in three steps. First, we use one language model to embed patent titles and abstracts and task descriptions in a shared semantic space, and a separate model to embed patent titles and abstracts and tool descriptions in another. Second, we compute pairwise patent–task and patent–tool cosine similarities. Third, we retain only the strongest matches using percentile-based similarity thresholds.

An embedding converts a piece of text into a vector in a high-dimensional semantic space designed to capture its contextual meaning. Using such embeddings therefore allows us to identify a close relationship between phrases such as "a method for classifying electronic documents" and "sort or classify information according to guidelines," even when they share relatively few words. We measure the similarity between two embeddings using cosine similarity, which increases as the vectors point in more similar directions.[3]

Because occupational tasks and tools differ in their textual content, we construct their embeddings using two domain-specific sentence-transformer models. For patent–task comparisons, which underlie automation exposure, we use JobBERT V2 (Decorte et al. 2025), an updated version of the JobBERT framework developed by Decorte et al. (2021). JobBERT V2 is trained on

[3]Formally, for texts represented by vectors $\mathbf{a}$ and $\mathbf{b}$, cosine similarity is given by $\cos(\mathbf{a}, \mathbf{b}) = \mathbf{a}^\top \mathbf{b} / (\|\mathbf{a}\| \|\mathbf{b}\|)$. Because the embeddings are normalized to unit length, this expression reduces to their dot product, $\mathbf{a}^\top \mathbf{b}$. Higher values indicate greater semantic similarity.

millions of job postings enriched with skill information and is therefore well suited to identifying semantic relationships between patent text and descriptions of the activities workers perform. For patent–tool comparisons, which underlie augmentation exposure, we use PatentSBERTa (Bekamiri et al. 2024), a sentence-transformer model trained specifically on patent documents. We choose this model because occupational tools refer primarily to technological artifacts and inputs—such as equipment, hardware, and instruments—that are conceptually closer to the technological objects described in patents than to the activities described in job postings.[4] We represent each patent by combining its title and abstract into a single text and, within each comparison, embed both sides of the match using the same model.

Based on these embeddings, we compute cosine similarity scores between every patent–task pair and every patent–tool pair. The resulting similarity matrices contain billions of pairs, the vast majority of which reflect weak semantic relationships. To isolate the strongest links, we convert the continuous similarity scores into binary matches using percentile thresholds from the corresponding pooled similarity distributions. Specifically, we define

$$H_{tp}^{\text{task}} = \mathbf{1}\left[s_{tp} \geq \tau^{\text{task}}\right], \qquad H_{rp}^{\text{tool}} = \mathbf{1}\left[s_{rp} \geq \tau^{\text{tool}}\right], \tag{11}$$

where s_{tp} and s_{rp} denote the cosine similarities between patent p and task t and between patent p and tool r, respectively. The indicators H_{tp}^{task} and H_{rp}^{tool} equal one when the corresponding similarity reaches the channel-specific threshold. In our baseline specification, τ^{task} and τ^{tool} are set at the 90th percentiles of the pooled patent–task and patent–tool similarity distributions. Appendix Figure C.1 shows the pooled distributions used to define these thresholds, together with the corresponding AI and non-AI distributions.

Figure 2 illustrates the patent–task and patent–tool matching procedure using one occupation-specific example from each channel. Panel A shows how an AI patent that automatically extracts and enters information from heterogeneous documents (US 10,817,656) matches tasks performed by statistical assistants, including entering, checking, and updating data. Panel B shows how an AI patent for intensity-modulated radiation treatment (US 7,180,980) matches tools used by radiation therapists, including medical linear accelerators and multileaf collimators.

Figure 3 broadens this illustration by comparing patent–task and patent–tool matches across representative occupations from all 22 two-digit SOC groups and representative patents from the eight AI technology functions listed in Appendix Table A.2.

[4] Although we deliberately use PatentSBERTa for patent–tool comparisons, the qualitative pattern of our main results is robust to constructing augmentation exposure with a general-purpose MPNet language model that is not trained specifically on patent text; see Appendix Table E.2.

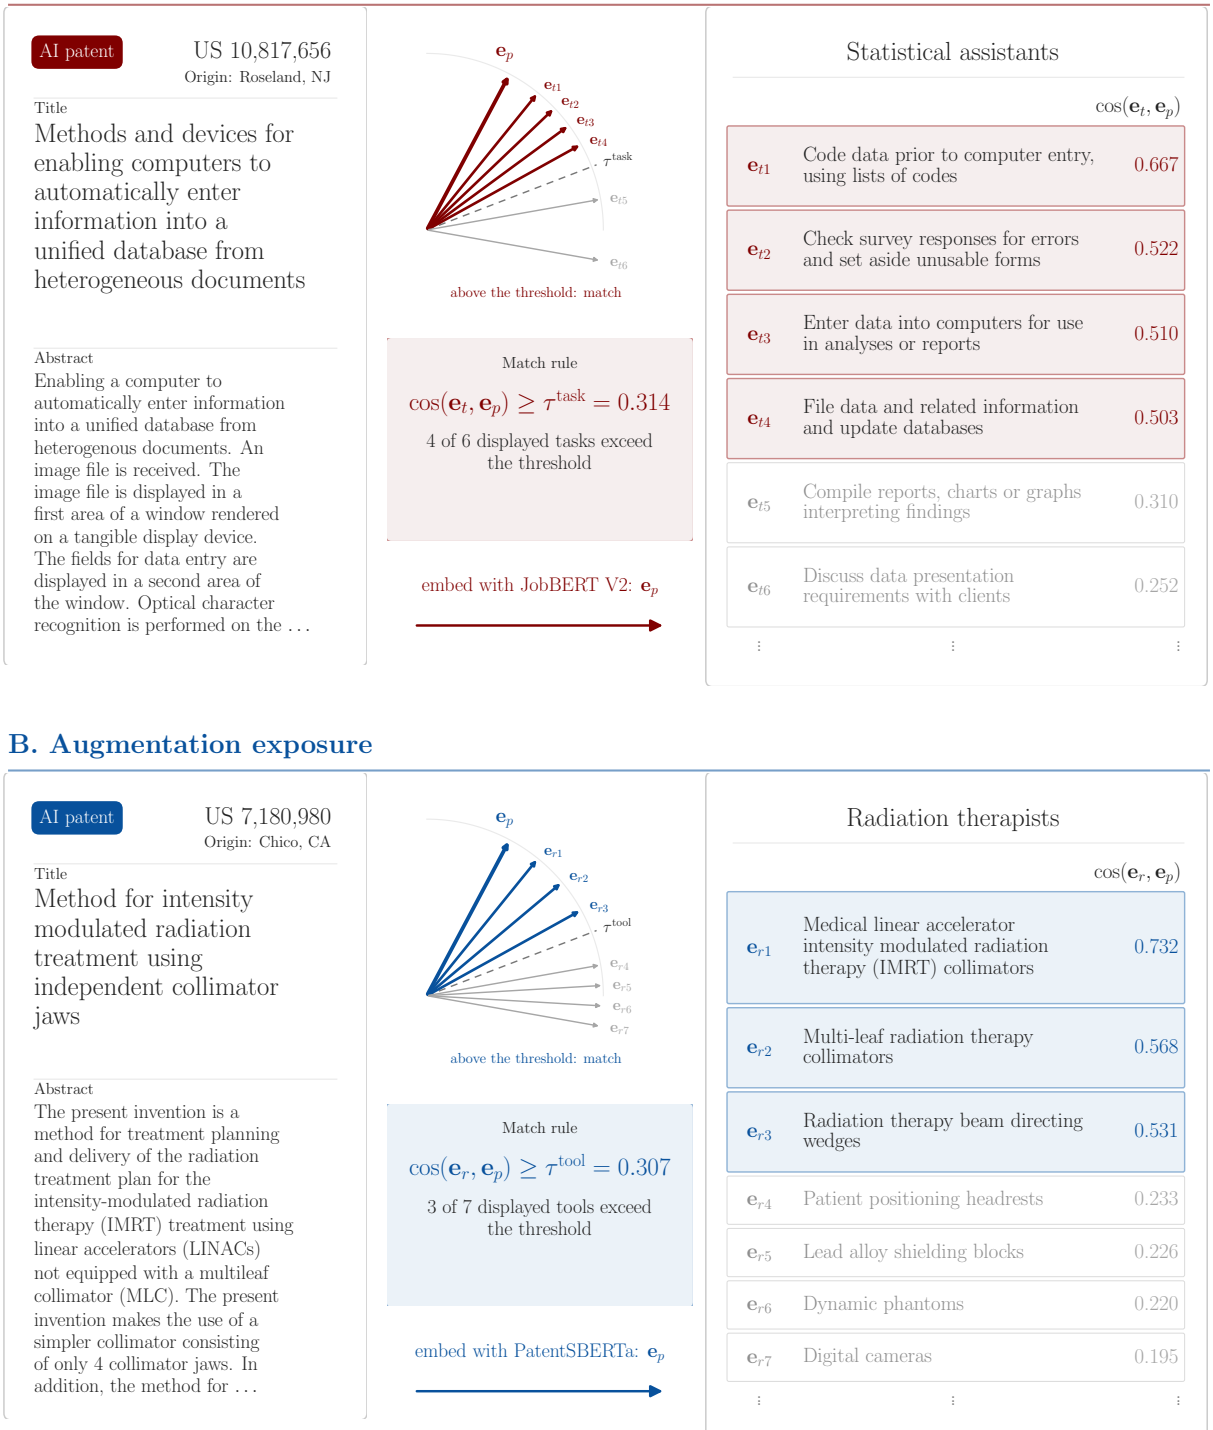


Figure 2: Occupation Examples of Patent–Task and Patent–Tool Matching

Notes: Panel A illustrates patent–task matching for statistical assistants using the JobBERT V2 language model, while Panel B illustrates patent–tool matching for radiation therapists using the PatentSBERTa language model. Colored entries meet the corresponding channel-specific 90th-percentile cosine-similarity threshold and are classified as matches. For readability, only six of the 14 tasks of statistical assistants and seven of the 63 tools of radiation therapists are displayed.

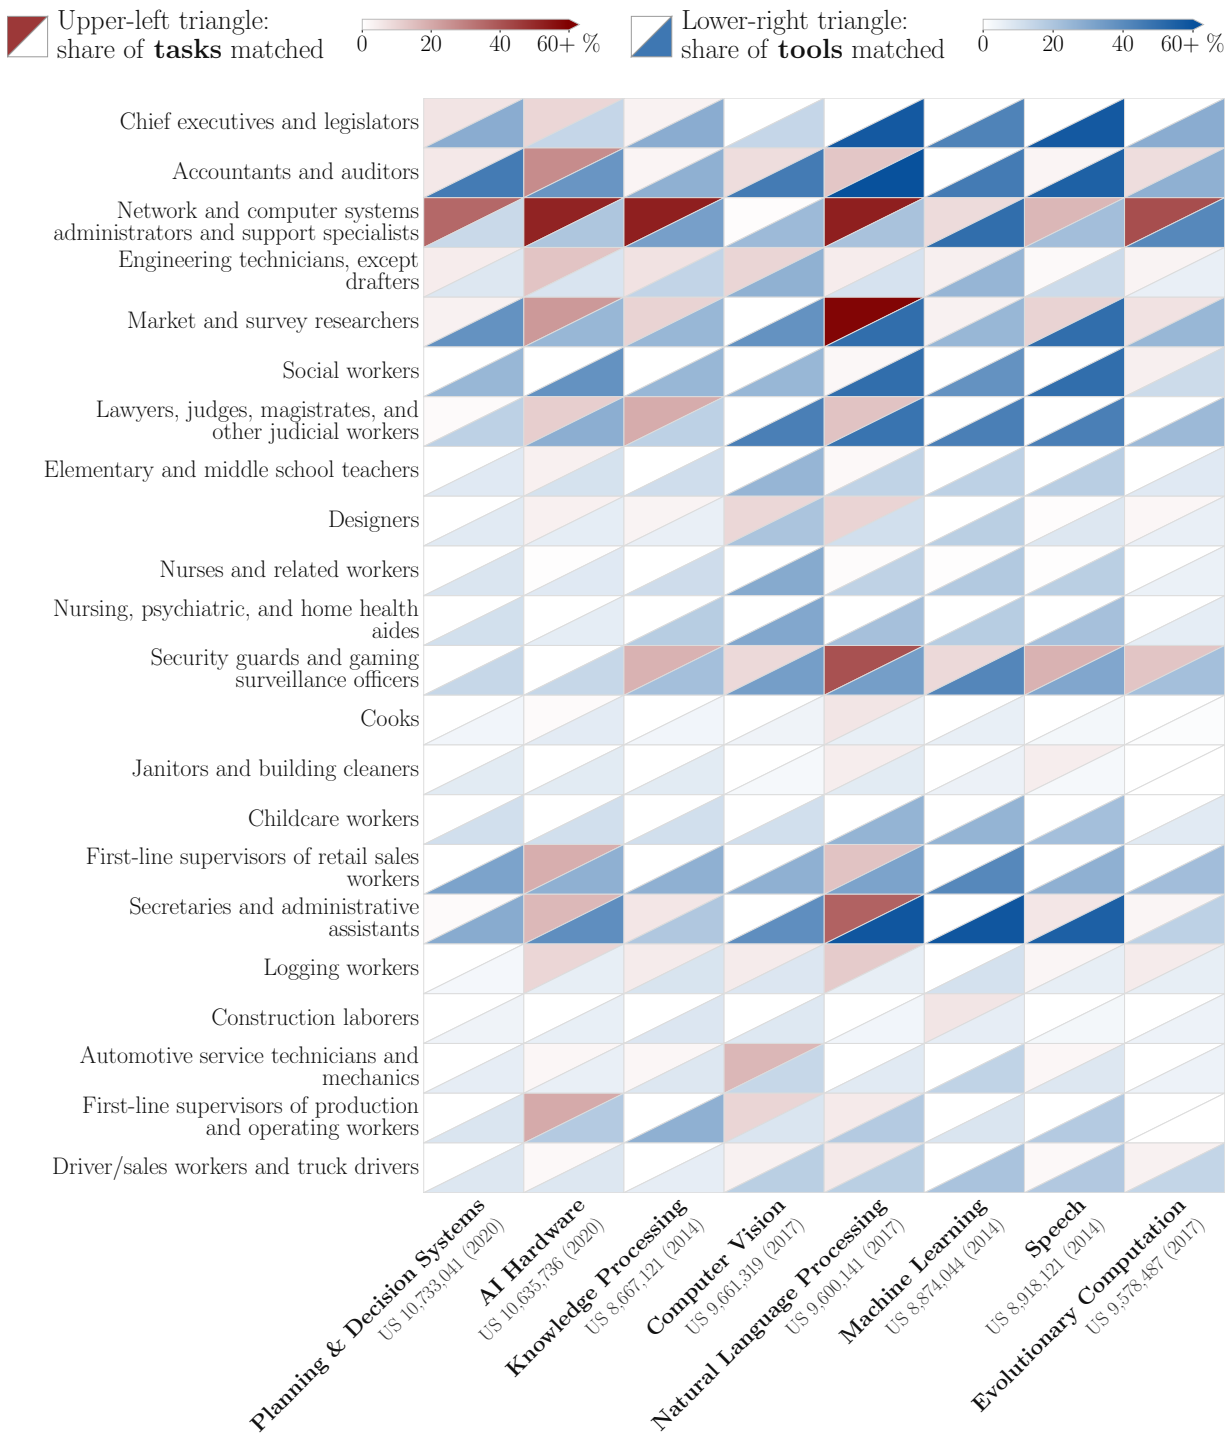


Figure 3: Patent–Task and Patent–Tool Matches Across Occupations and AI Functions

Notes: Rows show the detailed occupation with the largest national FTE employment share within each two-digit SOC group in 2005–07, excluding residual occupations ("other" or "n.e.c."). Columns show the eight AIPD technology functions. Each function is represented by the AI patent granted during 2005–2021 and assigned exclusively to that function whose title-and-abstract PatentSBERTa embedding has the highest cosine similarity to the mean embedding of the corresponding function-specific patent corpus. Cells report the unweighted shares of the occupation's O*NET tasks and tools satisfying the matching indicators in equation (11), using $\tau^{\text{task}} = 0.314$ and $\tau^{\text{tool}} = 0.307$. Linked USPTO patent records, in column order, are US 10,733,041, US 10,635,736, US 8,667,121, US 9,661,319, US 9,600,141, US 8,874,044, US 8,918,121, and US 9,578,487.

To construct Figure 3, we select the detailed occupation with the largest national full-time-equivalent (FTE) employment share within each group in the pooled 2005–07 American Community Survey (ACS) data. For each AI function, we select, among AI patents granted between 2005 and 2021, the patent whose embedding lies closest to the centroid of the function-specific patent corpus. An advantageous feature of our procedure is that patent–task and patent–tool matches are determined independently using the same underlying patent universe, so a patent is not assigned exclusively to either channel. The same innovation may therefore match occupational tasks, occupational tools, or both, allowing our measures to capture the possibility that it automates some activities while complementing workers in others.[5]

Because the same O*NET tools may be shared across a wide range of occupations, Figure 3 shows that tool matches are comparatively widespread across the grid. This pattern is particularly prominent among occupations whose O*NET inventories include general-purpose computing and communication tools. By contrast, task matches are more concentrated in specific occupation–technology pairs because O*NET tasks are defined at the occupation level.[6] Among the occupations displayed, task-match shares are especially high for network and computer systems administrators across several of the selected computing patents and for market and survey researchers under the information-retrieval patent (US 9,600,141), while the vehicle surround-view camera-calibration patent (US 9,661,319) produces its highest task-match share for automotive service technicians and mechanics. Because each AI function is represented by only one patent, Figure 3 is illustrative rather than a summary of the full distribution of AI innovation. Nevertheless, the occupation rankings based on average task- and tool-match shares across the eight displayed patents align closely with the corresponding national AI exposure measures, with rank correlations of 0.9 or above.

The next subsection aggregates these patent–task and patent–tool matches across all patents assigned to each commuting zone to construct the local occupation-level automation and augmentation exposure measures used in our main analysis.

3.3 Local Automation and Augmentation Exposure

We construct automation and augmentation exposure measures for each occupation–commuting-zone pair using a balanced panel of 430 occupations that can be consistently tracked across

[5]For a separate description of how the automation and augmentation composition of the U.S. patent stock evolves over time, we construct a patent-level classification based on each patent's maximum similarity to any occupational task or tool. This classification is exclusive by construction and is used only for the descriptive analysis in Appendix D.

[6]Appendix A.4 describes these O*NET task and tool data characteristics in more detail.

Census and ACS occupation classifications from 2000 onward.[7] Local labor markets are defined using the commuting zones of Tolbert and Sizer (1996), which group counties into economically integrated labor markets based on commuting patterns. Patents granted between 2005 and 2021 are assigned to commuting zones based on assignee locations and, when assignee locations are unavailable, inventor locations.[8] Let $\mathcal{P}_c$ denote the resulting set of patents assigned to commuting zone c.

Automation Exposure. Let $\mathcal{T}_o$ denote the set of tasks associated with occupation o. To account for differences in task importance, we use O*NET's importance scores w_{ot}, which we normalize within occupations so that the weights sum to one.

Using these weights, we define automation exposure for occupation o in commuting zone c as

$$\text{Aut}_{oc} = \sum_{t \in \mathcal{T}_o} \omega_{ot}^{\text{task}} \sum_{p \in \mathcal{P}_c} H_{tp}^{\text{task}}, \qquad \omega_{ot}^{\text{task}} = \frac{w_{ot}}{\sum_{t' \in \mathcal{T}_o} w_{ot'}}, \tag{12}$$

where $\omega_{ot}^{\text{task}}$ is the normalized O*NET importance weight for task t in occupation o, and H_{tp}^{task} is the patent–task match defined in equation (11). An occupation receives greater automation exposure when patents assigned to its commuting zone generate more matches with its tasks, with matches to more important tasks receiving greater weight.

Augmentation Exposure. Let $\mathcal{R}_o$ denote the set of tools associated with occupation o. Because O*NET does not provide importance scores for tools, we construct occupation-specific weights based on how closely each tool relates to the occupation's tasks. The intuition behind this approach is straightforward: workers use tools to perform tasks, so a tool that is semantically closely related to an occupation's overall task profile is likely to be more relevant to that occupation.

We first represent occupation o by an importance-weighted average of its task embeddings, normalized to unit length:

$$\bar{\mathbf{e}}_o^{\text{task}} = \frac{\sum_{t \in \mathcal{T}_o} \omega_{ot}^{\text{task}} \mathbf{e}_t^J}{\left\| \sum_{t \in \mathcal{T}_o} \omega_{ot}^{\text{task}} \mathbf{e}_t^J \right\|}, \tag{13}$$

[7] Of the 430 occupations, 416 are assigned directly to O*NET task and tool inventories. For the remaining 14 residual occupation categories without O*NET counterparts, we impute exposure using the average exposure of directly matched occupations in the corresponding broader occupational category. The construction of the occupation panel is described in Appendix A.2, while the assignment of O*NET tasks and tools to this panel is documented in Appendix A.4.

[8] Among geographically assigned domestic patents with assignee information, 98.4 percent have a single assignee, while approximately 14 percent of geographically assigned domestic patents rely on the inventor-location fallback. Further details on geographic assignment are provided in Appendix A.5.

where $\mathbf{e}_t^J$ denotes the JobBERT V2 embedding of task t, and $\bar{\mathbf{e}}_o^{\text{task}}$ denotes the unit-normalized importance-weighted task profile of occupation o. We assign each tool a raw relevance score equal to its cosine similarity with this task profile, set negative similarities to zero, and normalize the resulting scores within occupations:

$$\widetilde{w}_{or}^{\text{tool}} = \max\left\{0, \cos\left(\mathbf{e}_r^J, \bar{\mathbf{e}}_o^{\text{task}}\right)\right\}, \qquad \omega_{or}^{\text{tool}} = \frac{\widetilde{w}_{or}^{\text{tool}}}{\sum_{r' \in \mathcal{R}_o} \widetilde{w}_{or'}^{\text{tool}}}, \tag{14}$$

where $\mathbf{e}_r^J$ denotes the JobBERT V2 embedding of tool r, $\widetilde{w}_{or}^{\text{tool}}$ is its raw relevance score for occupation o, and $\omega_{or}^{\text{tool}}$ is the corresponding normalized tool weight. For this weighting step, tasks and tools are embedded in a common JobBERT V2 space so that their cosine similarity provides a meaningful comparison, while patent–tool matches remain those defined above using PatentSBERTa.[9]

Local augmentation exposure is then defined as

$$\text{Aug}_{oc} = \sum_{r \in \mathcal{R}_o} \omega_{or}^{\text{tool}} \sum_{p \in \mathcal{P}_c} H_{rp}^{\text{tool}}, \tag{15}$$

so that occupations receive greater augmentation exposure when patents assigned to their commuting zone generate more matches with their tools, with matches to more relevant tools receiving greater weight.

AI and Non-AI Exposure. We next use the AI classification described in Section 3.1 and decompose both automation and augmentation exposure into AI and non-AI components. Let D_p^{AI} equal one if patent p is classified as AI-related and zero otherwise. AI automation exposure is defined as

$$\text{Aut}_{oc}^{\text{AI}} = \sum_{t \in \mathcal{T}_o} \omega_{ot}^{\text{task}} \sum_{p \in \mathcal{P}_c} D_p^{\text{AI}} H_{tp}^{\text{task}}, \tag{16}$$

while AI augmentation exposure is defined as

$$\text{Aug}_{oc}^{\text{AI}} = \sum_{r \in \mathcal{R}_o} \omega_{or}^{\text{tool}} \sum_{p \in \mathcal{P}_c} D_p^{\text{AI}} H_{rp}^{\text{tool}}. \tag{17}$$

[9] To assess whether using task information to weight tools mechanically influences the relationship between our automation and augmentation measures, we construct alternative tool weights based on the similarity between tool descriptions and occupation-title profiles constructed from O*NET Alternate Titles. The resulting alternative weights are highly correlated with our preferred task-profile-based weights, and our results are robust to this alternative weighting procedure; see Appendix Table E.2.

The corresponding non-AI measures, denoted $\text{Aut}_{oc}^{\text{NonAI}}$ and $\text{Aug}_{oc}^{\text{NonAI}}$, are obtained by replacing D_p^{AI} with $(1 - D_p^{\text{AI}})$. By construction, the untransformed AI and non-AI measures sum to total exposure within each channel.

Relative AI Exposure. To distinguish the composition of technological exposure from its overall intensity, we construct relative measures equal to the AI share of total automation or augmentation exposure:

$$\text{Aut}_{oc}^{\text{RelAI}} = \frac{\text{Aut}_{oc}^{\text{AI}}}{\text{Aut}_{oc}^{\text{AI}} + \text{Aut}_{oc}^{\text{NonAI}}}, \qquad \text{Aug}_{oc}^{\text{RelAI}} = \frac{\text{Aug}_{oc}^{\text{AI}}}{\text{Aug}_{oc}^{\text{AI}} + \text{Aug}_{oc}^{\text{NonAI}}}. \tag{18}$$

By construction, these measures range from zero to one, with higher values indicating a larger AI share within the corresponding channel. Thus, two occupation–commuting-zone cells may face the same total amount of automation or augmentation exposure but differ in the share attributable to AI. Appendix Figures C.2 and C.3 show the spatial distributions of the average level and relative exposure measures, respectively, across the 741 U.S. commuting zones.

Transformation. The raw level-exposure measures are highly right-skewed and contain many zeros. We therefore apply the inverse hyperbolic sine (IHS) transformation to all level-exposure measures, which compresses the upper tail while retaining occupation–commuting-zone cells with zero exposure. The relative measures in equation (18) are calculated from the untransformed level measures and are not themselves transformed. Because these measures are defined only for cells with positive total exposure in the corresponding channel, cells with zero total exposure have undefined relative AI shares and are omitted from analyses involving the respective measure.

4 Descriptive Patterns of AI Exposure

Before turning to identification and the labor-market analysis, Figure 4 illustrates how AI automation and augmentation exposure vary across detailed occupations. Panel A reports employment-weighted occupation averages of the standardized AI exposure measures. Although the automation and augmentation measures are strongly related, occupations differ in their relative orientation toward the two channels. We illustrate this dispersion around the equality line by marking occupations with comparatively greater automation exposure in maroon and those with comparatively greater augmentation exposure in blue. Computer programmers and database administrators, for example, exhibit comparatively greater automation exposure, whereas childcare workers and some managerial and service occupations show greater augmentation exposure.

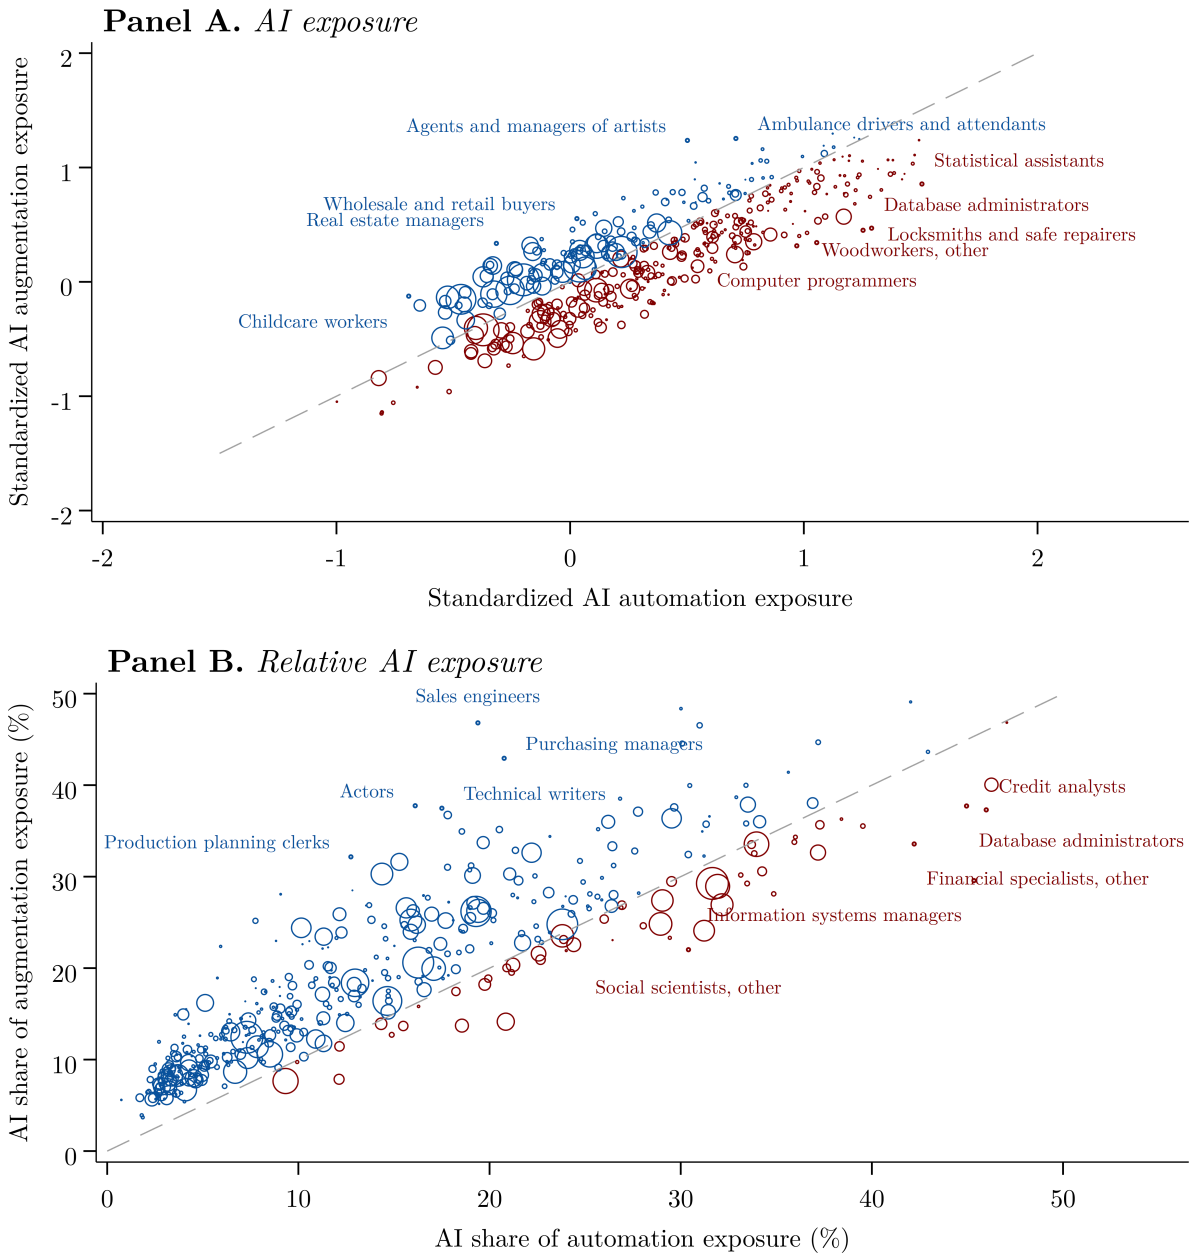


Figure 4: Occupation-Level AI Automation and Augmentation Exposure

Notes: Panel A reports employment-weighted occupation averages across commuting zones of the standardized inverse-hyperbolic-sine-transformed AI automation and augmentation exposure measures constructed from patents granted during 2005–2021. Panel B reports the corresponding occupation averages of the AI share of total exposure within each channel. The dashed line denotes equality between the measures shown on the two axes. Maroon observations lie below the equality line and blue observations lie above it, indicating relatively greater automation and augmentation exposure, respectively, in Panel A and a relatively larger AI share within the corresponding channel in Panel B. Occupation labels identify the five observations furthest from the equality line on each side of each panel. Bubble sizes are proportional to occupations' national FTE employment shares in 2005–07. The figure includes all 416 occupations directly assigned to O*NET task and tool inventories and excludes the 14 residual n.e.c. occupations with imputed exposure measures.

Panel B plots the AI share of total automation exposure against the AI share of total augmentation exposure, using the same colors to indicate whether AI accounts for a larger share of exposure within the automation or augmentation channel. For most occupations, AI represents a larger share of augmentation exposure than of automation exposure, although the size of this difference varies considerably across occupations. Sales engineers, purchasing managers, and actors exhibit substantially larger AI shares within augmentation than automation, whereas computer and information systems managers and several financial occupations exhibit the opposite pattern. Appendix Table C.1 provides complementary evidence by listing the ten occupations most highly exposed under each measure.

These occupation-level averages, however, conceal considerable geographic variation. To demonstrate this, Table 1 decomposes the employment-weighted variation in each exposure measure into variation across commuting zones within detailed occupations and variation between occupation means. For AI automation and augmentation exposure, 85.3 percent of the variation in automation and 87.9 percent of the variation in augmentation occurs within occupations across commuting zones. Hence, workers in the same detailed occupation may face substantially different exposure depending on the innovative activity associated with their local labor market.

Geographic variation also remains important when considering the composition rather than the overall intensity of technological exposure. Within-occupation variation accounts for 57.9 percent of the variation in relative AI automation exposure and 56.1 percent of the corresponding variation in relative AI augmentation exposure. Assigning a single national exposure measure to each occupation would therefore neglect a substantial share of the variation in both the intensity and AI composition of automating and augmenting technological change. This local variation motivates the spatial identification strategy introduced in the next section.

Table 1: Within- and Between-Occupation Variation in AI Exposure

	Descriptive statistics		Variance decomposition (%)		Sample		
Exposure measure	Mean	SD	Within	Between	Observations	Occupations	CZs
AI automation	4.267	2.665	85.3	14.7	91,595	416	741
AI augmentation	5.165	2.781	87.9	12.1	91,595	416	741
Relative AI automation	0.145	0.142	57.9	42.1	89,183	416	720
Relative AI augmentation	0.190	0.135	56.1	43.9	89,206	416	720

Notes: The table reports employment-weighted descriptive statistics and variance decompositions across occupation–commuting-zone cells. The within component captures variation across commuting zones within detailed occupations, while the between component captures variation in employment-weighted occupation means. The two components are reported as percentages of total variance and sum to 100 within each row. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed but not standardized. Relative AI exposure is the AI share of total exposure within the corresponding channel and is defined only for cells with positive total exposure in that channel. The sample includes 416 occupations and excludes the 14 residual n.e.c. occupations with imputed exposure measures.

5 Identification Strategy

The key challenge in estimating the labor-market effects of local AI exposure is that this exposure is not randomly assigned. Regions with stronger employment growth, higher productivity, or more favorable economic conditions may both generate AI-related innovations and adopt AI technologies more rapidly. Local industrial specialization and workforce characteristics may similarly influence both the development and adoption of AI technologies and subsequent labor-market outcomes. Observed relationships between local AI exposure and changes in employment or wages may therefore reflect underlying economic trends rather than the causal effects of local AI exposure. This concern is particularly relevant in our setting because AI-related automating and augmenting innovation is concentrated in economically dynamic areas, such as Silicon Valley and the Northeast Corridor, and may extend into nearby labor markets through geographically localized knowledge and technology spillovers. To address these concerns, we develop an instrumental-variable strategy that exploits the historical geography of breakthrough AI innovations and their subsequent spatial diffusion across U.S. local labor markets. In this strategy, spatial diffusion refers to the geographic propagation of innovative activity, as reflected in subsequent local patenting, rather than to directly observed workplace adoption.

5.1 Spatial Diffusion Instrument

We first identify historical breakthrough patents granted between 1980 and 2000 and distinguish which of these patents are AI-related.[10] To do so, we combine the text-based breakthrough-patent measure developed by Kelly et al. (2021) with the AI patent classification of Pairolero et al. (2025). Kelly et al. (2021) characterize breakthrough patents as inventions that are novel relative to prior innovation but closely related to subsequent technological developments.[11] AI patenting was particularly breakthrough-intensive during this historical window: 59 percent of AI patents granted between 1980 and 2000 satisfy our breakthrough criterion, compared with 14 percent of non-AI patents. Further details on the patent data and breakthrough construction are provided in Appendix A.5. These historical AI breakthroughs substantially predate the recent period

[10] We use 1980–2000 as our baseline historical window because it provides a broad pre-period for capturing the geography of frontier AI innovation while ending before the 2005–2021 period used to construct contemporary exposure. As a stricter temporal robustness check, Appendix Table H.1 reconstructs the instruments using only breakthrough patents granted during 1980–1990. The resulting estimates preserve the qualitative pattern of the baseline results, although the estimated effects are larger in magnitude.

[11] The measure assigns greater importance to patents that combine low textual similarity to preceding patents with high similarity to patents granted over the subsequent ten years. Autor et al. (2024) use lagged breakthrough patents as instruments for subsequent flows of automating and augmenting innovation. Following their implementation, we classify a patent as a breakthrough if its importance lies in the top decile of the corresponding grant-year distribution.

of widespread AI diffusion and adoption. Building on evidence that new technologies spread gradually from the places where they first emerge (Kalyani et al. 2025), we use the locations of these historical breakthroughs as the starting points for constructing spatial exposure.[12]

Because the patent–task and patent–tool matching procedure described in Section 3 is applied to the entire patent universe from 1980 to 2021, we can characterize historical AI breakthroughs according to their automation and augmentation relevance for each occupation. To define spatial proximity, we represent commuting-zone geography as a contiguity network in which adjacent commuting zones are connected and define rings around each focal commuting zone by shortest-path distance in this network. For every occupation–commuting-zone pair, we aggregate occupation-specific breakthrough exposure across commuting zones of origin using inverse-ring-distance weights. Instrumental exposure is therefore greater when historical AI breakthroughs are more relevant to an occupation's tasks or tools and originate closer to the focal commuting zone in the contiguity network.

Our baseline instrument excludes breakthrough patents originating in the focal commuting zone and the first two surrounding contiguity rings.[13] Imposing this buffer reduces the concern that historical breakthroughs reflect persistent local inventive capacity or shared economic conditions within nearby innovation networks that could independently affect subsequent labor-market outcomes. Appendix Figure F.2 provides descriptive support for this choice by showing that contemporaneous spatial correlation in AI exposure is concentrated at short contiguity-ring distances and declines sharply with distance within both channels. Appendix Table F.1 shows that historical AI breakthrough exposure continues to predict subsequent local AI exposure far beyond this buffer, with its predictive content generally declining with geographic distance. These patterns play important but distinct roles in our design: the former motivates the exclusion buffer, whereas the latter documents the geographic gradient in instrument relevance underlying our distance-weighting approach.

For occupation o in commuting zone c, the instrument for channel $k \in \{\text{Aut}, \text{Aug}\}$ is

$$Z_{oc}^{k,\text{AI}} = \sum_{j:\, d_{cj} \geq 3} w_{cj} BT_{oj}^{k,\text{AI}}, \qquad w_{cj} = \frac{1/d_{cj}}{\sum_{\ell:\, d_{c\ell} \geq 3} 1/d_{c\ell}}. \tag{19}$$

[12] We use the same USPTO location information and assignment procedure as for our contemporaneous exposure measures. Patents assigned to locations outside the United States or for which no location can be determined are excluded; see Appendix A.5 for details.

[13] Seven geographically isolated commuting zones—one in Alaska, four in Hawaii, Nantucket County in Massachusetts, and San Juan County in Washington—have no eligible donors after imposing the two-ring exclusion buffer and are assigned zero instrumental exposure. Other Alaskan commuting zones retain eligible donors within their connected networks. Excluding the seven affected commuting zones leaves our results virtually unchanged.

where $BT_{oj}^{k,\mathrm{AI}}$ denotes occupation-specific exposure in channel k to historical AI breakthrough patents originating in commuting zone j. The variable d_{cj} denotes the shortest-path distance between commuting zones c and j in the contiguity network, measured in rings. Thus, $d_{cj} \geq 3$ excludes the focal commuting zone and its first two surrounding rings. For focal commuting zones with eligible donors, the inverse-ring-distance weights w_{cj} are normalized to sum to one across all eligible origin commuting zones. Figure 5 illustrates the exclusion buffer and inverse-ring-distance weighting used to construct the instrument.

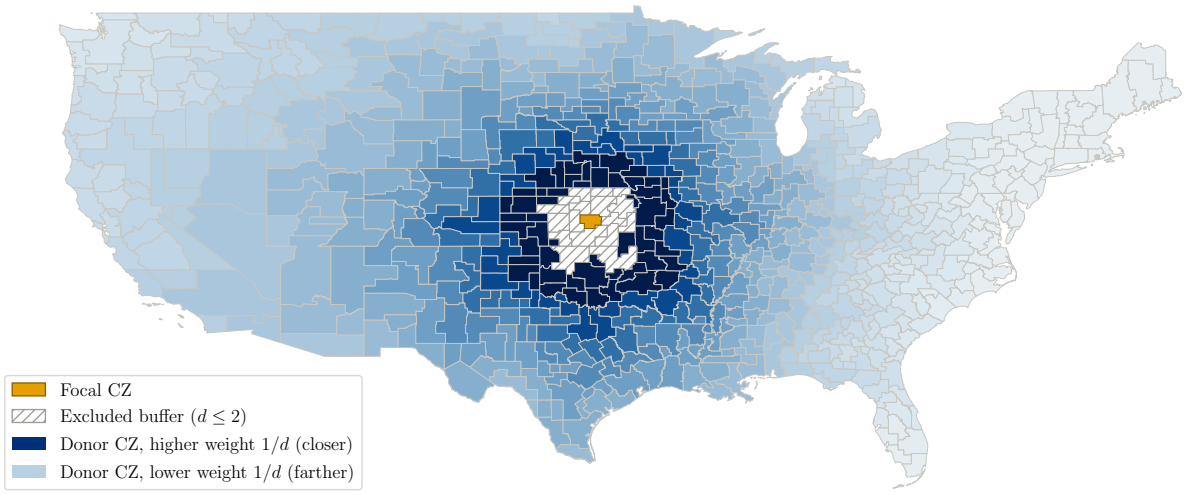


Figure 5: Spatial Construction of the Historical AI Breakthrough Instruments

Notes: The figure illustrates the spatial weighting used to construct the historical instruments from AI breakthrough patents granted during 1980–2000. Breakthroughs originating in the focal commuting zone or its first two contiguity rings are excluded. Eligible donor commuting zones at contiguity distance $d \geq 3$ receive weights proportional to $1/d$, normalized to sum to one across eligible donors, where d is the shortest-path distance in the commuting-zone contiguity network. Alaska and Hawaii are omitted from the map for visual clarity but remain in the sample.

5.2 Econometric Specification

Using the spatial diffusion instruments developed above, we estimate the following second-stage specification:

$$\Delta Y_{oc} = \beta_1 \mathrm{Aut}_{oc}^{\mathrm{AI}} + \beta_2 \mathrm{Aug}_{oc}^{\mathrm{AI}} + \delta_1 \mathrm{Aut}_{oc}^{\mathrm{NonAI}} + \delta_2 \mathrm{Aug}_{oc}^{\mathrm{NonAI}} + \mathbf{X}_{oc}'\boldsymbol{\gamma} + \lambda_{c,b(o)} + \varepsilon_{oc}, \tag{20}$$

where ΔY_{oc} denotes the long-run change in a labor-market outcome for occupation o in commuting zone c. We measure local occupational labor-market outcomes using the ACS, pooling observations over 2005–07 and 2022–24. We restrict our sample to occupation–commuting-zone cells with at least ten observations in each period. Outcomes are the log change in FTE employment, the log change in the wage bill, and the change in mean log hourly wage between the

two periods. Further details on the sample restrictions and variable construction are provided in Appendix A.1.

The variables $\text{Aut}_{oc}^{\text{AI}}$ and $\text{Aug}_{oc}^{\text{AI}}$ measure local occupational exposure to AI-driven automation and augmentation, respectively. Their non-AI counterparts, $\text{Aut}_{oc}^{\text{NonAI}}$ and $\text{Aug}_{oc}^{\text{NonAI}}$, are constructed analogously using patents not classified as AI-related. The vector $\mathbf{X}_{oc}$ contains baseline controls, including the shares of workers below age 30, above age 55, female, foreign-born, and employed in manufacturing. In an augmented specification, we additionally control for baseline college share and occupational routine intensity.[14] The term $\lambda_{c,b(o)}$ denotes commuting-zone-by-broad-occupation fixed effects, with occupations classified into five broad groups and $b(o)$ indicating the group to which occupation o belongs.[15]

The two endogenous AI exposure measures in equation (20) are jointly instrumented using the spatial exposure measures constructed from historical AI breakthroughs:

$$\begin{aligned} \text{Aut}_{oc}^{\text{AI}} &= \pi_{11} Z_{oc}^{\text{Aut,AI}} + \pi_{12} Z_{oc}^{\text{Aug,AI}} + \mathbf{W}_{oc}'\boldsymbol{\theta}_1 + \lambda_{c,b(o)} + u_{oc}, \\ \text{Aug}_{oc}^{\text{AI}} &= \pi_{21} Z_{oc}^{\text{Aut,AI}} + \pi_{22} Z_{oc}^{\text{Aug,AI}} + \mathbf{W}_{oc}'\boldsymbol{\theta}_2 + \lambda_{c,b(o)} + v_{oc}, \end{aligned} \tag{21}$$

where $\mathbf{W}_{oc} = \left(\text{Aut}_{oc}^{\text{NonAI}}, \text{Aug}_{oc}^{\text{NonAI}}, \mathbf{X}_{oc}'\right)'$ collects the included non-AI exposure measures and baseline controls.

To examine whether the composition of technological change matters in addition to its overall intensity, we estimate

$$\Delta Y_{oc} = \beta_1 \text{Aut}_{oc}^{\text{RelAI}} + \beta_2 \text{Aug}_{oc}^{\text{RelAI}} + \delta_1 \text{Aut}_{oc} + \delta_2 \text{Aug}_{oc} + \mathbf{X}_{oc}'\boldsymbol{\gamma} + \lambda_{c,b(o)} + \varepsilon_{oc}, \tag{22}$$

where the variables $\text{Aut}_{oc}^{\text{RelAI}}$ and $\text{Aug}_{oc}^{\text{RelAI}}$ denote the AI shares of total automation and augmentation exposure defined in equation (18). Conditional on total exposure within each channel (Aut_{oc} and Aug_{oc}), β_1 and β_2 capture the effect of shifting the composition of automating and augmenting innovation toward AI and away from non-AI innovation.

For each channel $k \in \{\text{Aut}, \text{Aug}\}$, we apply the same spatial weighting procedure to AI and non-AI historical breakthroughs and, whenever total historical breakthrough exposure in channel k is positive, construct the relative instrument as

[14] Following Firpo et al. (2011), we construct a routine-intensity index combining five O*NET Work Context measures: degree of automation, importance of repeating the same tasks, structured versus unstructured work, pace determined by the speed of equipment, and time spent making repetitive motions. Appendix A.3 describes the construction in detail.

[15] The five broad occupational groups follow the broad groupings of the 2010 Census occupational classification and aggregate the following two-digit SOC major groups: management, business, science, and arts occupations (11–29); service occupations (31–39); sales and office occupations (41–43); natural resources, construction, and maintenance occupations (45–49); and production, transportation, and material moving occupations (51–53).

$$Z_{oc}^{k,\text{RelAI}} = \frac{Z_{oc}^{k,\text{AI}}}{Z_{oc}^{k,\text{AI}} + Z_{oc}^{k,\text{NonAI}}}. \tag{23}$$

The two relative AI exposure measures are jointly instrumented using their corresponding relative breakthrough-exposure measures. The corresponding first stages include total automation and augmentation exposure, the baseline controls, and the same fixed effects as equation (22). All specifications are weighted by 2005–07 FTE employment shares, and standard errors are clustered at the commuting-zone level.

5.3 Instrument Relevance and Identifying Assumptions

A second identification challenge in our setting is separately identifying the labor-market effects of AI-driven automation and augmentation. Although conceptually distinct, both channels may arise from the same locally generated innovations, resulting in a high correlation between the exposure measures across occupation–commuting-zone pairs (corr = 0.95). Identification therefore requires instruments that generate sufficiently distinct variation across the two channels. Historical AI breakthroughs are particularly useful in this regard because they represent novel technological advances that reveal early, occupation-specific links between AI technologies and workers' tasks and tools. Reassuringly, we find that the corresponding instruments are substantially less correlated (corr = 0.48), as illustrated in Appendix Figure F.3.

More formally, Table 2 shows that the automation instrument has a substantially larger loading in the automation first stage than in the augmentation first stage, while the augmentation instrument loads more strongly on augmentation than on automation. Although some cross-loading is expected because our agnostic measurement strategy described in Section 3 allows the same innovation to contribute to both channels through independent patent–task and patent–tool matches, the larger own-channel loadings indicate that the instruments capture distinct dimensions of AI-related technological change. The Sanderson–Windmeijer statistics indicate strong conditional first stages for both endogenous exposure measures, while the Kleibergen–Paap F-statistic indicates strong joint relevance of the excluded instruments. A similar pattern holds for the relative AI specifications.

The validity of our identification strategy requires that, conditional on the controls and fixed effects included in the specification, spatial exposure to historical AI breakthroughs does not affect post-baseline changes in local occupational labor-market outcomes through channels other than contemporaneous local AI automation and augmentation exposure. Several features of the empirical design mitigate potential threats to this assumption. First, following the approach of

Table 2: First-Stage Regressions for AI and Relative AI Exposure

	AI exposure		Relative AI exposure	
	(1) Automation	(2) Augmentation	(3) Automation	(4) Augmentation
Automation IV	0.327***	-0.002		
	(0.009)	(0.002)		
Augmentation IV	0.117***	0.355***		
	(0.007)	(0.006)		
Relative automation IV			8.603***	0.372***
			(0.304)	(0.046)
Relative augmentation IV			1.467***	11.310***
			(0.276)	(0.296)
Baseline workforce controls	✓	✓	✓	✓
Non-AI exposure controls	✓	✓		
Total exposure controls			✓	✓
CZ × broad occupation FE	✓	✓	✓	✓
SW F-statistic	1494.65	3575.49	1412.10	1757.03
KP F-statistic	711.01		463.04	
Observations	93,026	93,026	89,001	89,001

Notes: The table reports first-stage regressions from the preferred 2SLS specifications. Columns (1)–(2) use standardized AI automation and augmentation exposure, while columns (3)–(4) use standardized relative AI exposure, defined as the AI share of total exposure within each channel. The relative specifications use the common sample for which both relative exposure measures and instruments are defined, requiring positive total local and historical breakthrough exposure in both channels. Excluded instruments are distance-weighted AI breakthrough exposure constructed outside the two-ring exclusion buffer; relative instruments are the corresponding AI shares. All specifications are weighted by 2005–07 FTE employment shares. Sanderson–Windmeijer statistics assess conditional first-stage strength for each endogenous exposure, while the Kleibergen–Paap F-statistic summarizes the joint strength of the excluded instruments. Standard errors are clustered by commuting zone. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Autor et al. (2024), we construct the instruments exclusively from breakthrough patents rather than from the broader flow of AI patenting activity. These AI breakthrough patents were granted between 1980 and 2000 and therefore predate the outcome period and substantially predate the recent widespread adoption of AI technologies. Second, excluding breakthroughs originating in the focal commuting zone and its first two surrounding rings reduces the influence of persistent local and neighboring economic conditions associated with historical innovative activity. Third, we condition on baseline workforce characteristics, commuting-zone-by-broad-occupation fixed effects, and technological-exposure controls—non-AI automation and augmentation exposure in the baseline specifications and total exposure within each channel in the relative AI specifications—to account for broader patterns of automating and augmenting technological change.

We complement our design features with a broad set of robustness checks that vary the instrument construction, including larger exclusion buffers, alternative distance-weighting schemes, more narrowly defined sets of donor rings, an earlier breakthrough window, and the exclusion of major AI innovation hubs from the donor set, as documented in Appendices G and H.

6 Main Results

6.1 Effects of AI Automation and Augmentation

Table 3 presents the baseline 2SLS estimates of the effects of AI automation and augmentation exposure on three labor-market outcomes. The unit of observation is an occupation–commuting-zone cell, and the dependent variables are the log changes in the wage bill and FTE employment and the change in mean log hourly wage between 2005–07 and 2022–24. The IHS-transformed AI exposure measures are standardized within the common estimation sample, so that the reported coefficients capture the effect of a one-standard-deviation increase in exposure on the corresponding outcome. Columns (1)–(4) progressively introduce more demanding geographic and occupational fixed effects. Our preferred specification is reported in column (4), while column (5) additionally controls for baseline college share and routine intensity.[16]

The results reveal a clear contrast between AI automation and augmentation. Across all specifications, greater AI automation exposure is estimated to substantially reduce both wage bills and employment. In our preferred specification, a one-standard-deviation increase in AI automation exposure implies a reduction in the wage bill of 0.248 log points, equivalent to approximately 22 percent, and a reduction in employment of 0.231 log points (21 percent). The estimated effect on hourly wages is small and statistically insignificant. By contrast, a one-standard-deviation increase in AI augmentation exposure implies a 0.145-log-point rise in the wage bill (16 percent), reflecting gains of 0.079 log points in hourly wages (8 percent) and 0.067 log points in employment (7 percent). The estimated employment effect becomes smaller and statistically insignificant after controlling for baseline college share and occupational routine intensity in column (5), whereas the estimated wage effect is more robust across specifications, remaining positive and statistically significant throughout.

The finding that the decline in wage bills from AI automation primarily reflects employment losses rather than wage adjustments is consistent with the displacement effect of automation dominating any countervailing scale response, as described in Section 2. The corresponding hourly-wage response is less clear: the estimated effect is small, statistically insignificant in our preferred specification, and varies in sign across specifications. The muted average wage response may reflect heterogeneous effects across occupations that partly offset one another. This interpretation is consistent with recent evidence from Autor and Thompson (2025), which

[16] We treat baseline college share and routine intensity separately from the other controls because they capture dimensions of skill composition and task structure along which technological change may affect occupational labor demand heterogeneously. We consequently treat column (5) as a sensitivity check.

Table 3: Effects of AI Automation and Augmentation on Local Labor Markets

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.154***	-0.155***	-0.248***	-0.248***	-0.099***
	(0.033)	(0.028)	(0.021)	(0.019)	(0.021)
AI Augmentation	0.165***	0.223***	0.163***	0.145***	0.091***
	(0.048)	(0.047)	(0.027)	(0.026)	(0.024)
Panel B: Employment change					
AI Automation	-0.151***	-0.139***	-0.230***	-0.231***	-0.132***
	(0.028)	(0.024)	(0.018)	(0.016)	(0.018)
AI Augmentation	0.133***	0.147***	0.084***	0.067***	0.030
	(0.039)	(0.038)	(0.021)	(0.020)	(0.019)
Panel C: Hourly wage change					
AI Automation	-0.009	-0.018**	-0.013	-0.011	0.014*
	(0.008)	(0.007)	(0.008)	(0.007)	(0.008)
AI Augmentation	0.035***	0.068***	0.082***	0.079***	0.071***
	(0.013)	(0.012)	(0.010)	(0.010)	(0.010)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	486.04	623.53	551.09	711.01	719.79
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using the common estimation sample of column (5). The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages between 2005–07 and 2022–24. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this sample. They are instrumented using the corresponding distance-weighted exposure measures constructed from AI breakthrough patents granted during 1980–2000 outside the two-ring exclusion buffer. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

emphasizes that automation may raise or lower wages depending on how it changes the expertise required for tasks that workers continue to perform.

The results for AI augmentation point to a different pattern of adjustment. Whereas Autor et al. (2024) find that their output-based measure of augmentation predicts wage-bill growth almost entirely through employment gains, our tool-based measure yields a substantial positive effect on average hourly wages. This wage response is consistent with the productivity-enhancing role of tool augmentation emphasized in Section 2, although our estimates do not identify productivity directly.

The qualitative contrast between AI automation and augmentation, most clearly along the automation–employment and augmentation–wage margins, remains evident across a range of

measurement choices and instrument constructions.[17] These include alternative patent–task and patent–tool matching thresholds, job-title-profile-based tool weights instead of task-profile-based weights, the use of a general-purpose MPNet language model, and the alternative instrument constructions outlined in Section 5. The corresponding results and additional discussion are reported in the Appendix.

6.2 Effects Across Time Periods

The baseline estimates average technological exposure and labor-market adjustment over a period in which AI capabilities and innovative activity changed substantially. To examine whether the effects of automation and augmentation differ across this period, we split contemporary exposure into an early window, 2005–2013, and a late window, 2014–2021, and relate these measures to labor-market changes from 2005–07 to 2014–16 and from 2014–16 to 2022–24, respectively. Because the two outcome windows differ in length, outcome changes are decadalized, while the AI exposure measures are standardized within each period. Table 4 reports the resulting 2SLS estimates together with p-values for tests of whether the early- and late-period coefficients differ.

The results show that the contrast between the effects of automation and augmentation becomes clearly more pronounced in the later period. AI automation reduces employment in both periods, but the negative effect increases in magnitude from −0.098 to −0.158 log points, while the negative wage-bill effect more than doubles from −0.078 to −0.195 log points. The hourly-wage response also changes markedly, from 0.031 log points in the early period to −0.042 log points in the later period. This sign reversal indicates that the wage consequences of AI automation became substantially less favorable in the later period.

For augmentation, the pattern moves in the opposite direction. In the early period, AI augmentation raises hourly wages by 0.029 log points but has little effect on employment or the wage bill. In the later period, by contrast, it raises employment by 0.089 log points and hourly wages by 0.061 log points, producing a substantially larger wage-bill effect of 0.141 log points. The stronger estimated augmentation effects in the later period point to AI-driven augmentation as a potentially important countervailing force to the employment losses associated with AI automation. These temporal differences are supported by the equality tests, which reject equality of the early- and late-period coefficients for the wage-bill and employment effects of both channels and for the hourly-wage effect of automation, while the difference in augmentation's hourly-wage

[17] We additionally examine the sensitivity of statistical inference to alternative forms of residual dependence. Conley spatial-HAC standard errors are similar to the baseline commuting-zone-clustered standard errors, while statistical precision is substantially reduced under more conservative two-way clustering by commuting zone and detailed occupation; see Appendix Table F.3.

Table 4: Effects of AI Automation and Augmentation Across Time Periods

	Early period (1)	**Late period (2)**	**Equality test** p-value
Exposure window	2005–2013	2014–2021	
Outcome window	2005–07 to 2014–16	2014–16 to 2022–24	
Panel A: Decadalized wage bill change			
AI Automation	-0.078***	-0.195***	< 0.001
	(0.016)	(0.020)	
AI Augmentation	0.029	0.141***	0.002
	(0.021)	(0.025)	
Panel B: Decadalized employment change			
AI Automation	-0.098***	-0.158***	0.029
	(0.015)	(0.018)	
AI Augmentation	-0.007	0.089***	0.002
	(0.018)	(0.020)	
Panel C: Decadalized hourly wage change			
AI Automation	0.031***	-0.042***	< 0.001
	(0.006)	(0.008)	
AI Augmentation	0.029***	0.061***	0.058
	(0.009)	(0.012)	
KP F-statistic	492.95	527.45	
Observations	86,510	86,510	

Notes: Columns (1) and (2) report separate 2SLS estimates for the early and late exposure and outcome windows shown above. The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages. Outcomes are decadalized by multiplying the early-period changes by 10/9 and the late-period changes by 10/8. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized separately within each period. The exposure measures are instrumented using the corresponding historical-breakthrough instruments constructed from AI breakthrough patents granted during 1980–2000 outside the two-ring exclusion buffer. Both specifications control for period-specific non-AI automation and augmentation exposure and starting-period workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), include commuting-zone-by-broad-occupation fixed effects, and use the same balanced sample and 2005–07 FTE employment-share weights. Standard errors clustered by commuting zone are reported in parentheses. The final column reports p-values from stacked tests of equality between the early- and late-period coefficients, allowing instruments, controls, and fixed effects to vary by period. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

effect is significant only at the 10 percent level.

6.3 Relative AI Automation and Augmentation Effects

We next ask whether the technological composition within the two channels matters. In other words, does shifting automation or augmentation exposure toward AI and away from non-AI innovation alter the labor-market effects associated with those channels? This question is increasingly important because AI has become a rapidly growing component of technological change, as illustrated in Appendix Figure F.1. To answer it, we estimate equation (22) using the relative

AI exposure measures constructed in Section 3.3, which capture the AI share of total exposure within each channel, while controlling for total automation and augmentation exposure.

The preferred 2SLS estimates in column (4) of Table 5 suggest that increasing the AI share of automation exposure has adverse effects on both wage-bill and employment growth. A one-standard-deviation increase in relative AI automation exposure is associated with reductions of 0.030 log points in the wage bill, equivalent to approximately 3.0 percent, and 0.027 log points in employment (2.7 percent). The estimated effect on hourly wages is considerably smaller but positive, at 0.005 log points (0.5 percent). When baseline college share and routine intensity are added in column (5), the negative employment and wage-bill effects are substantially attenuated, whereas the positive hourly-wage effect becomes larger and more precisely estimated.

For augmentation, the positive but small and statistically insignificant wage-bill effect in column (4) conceals sizable and precisely estimated adjustments along the employment and wage margins. Conditional on overall augmentation exposure, a one-standard-deviation increase in relative AI augmentation exposure reduces employment by 0.027 log points (2.7 percent) while simultaneously raising hourly wages by 0.030 log points (3.0 percent). These opposing responses largely offset one another in the wage bill.

The findings presented in Table 5 therefore suggest that AI sharpens distinct dimensions of the labor-market consequences of the two channels, strengthening the employment-displacing effect of automation while shifting the effects of augmentation away from employment growth and toward wage gains. Direct comparisons of AI and non-AI level-exposure measures yield similar qualitative patterns, as shown in Appendix Table E.3. Taken together, this wage-oriented shift through the augmentation channel combined with stronger employment displacement through AI-related automation echoes recent concerns that AI may reinforce the labor-displacing direction of technological change in the twenty-first century (Acemoglu and Restrepo 2020; Acemoglu 2021; Acemoglu et al. 2026).

Finally, we examine whether the effects of relative AI exposure vary systematically with the skill composition and task structure of occupation–commuting-zone cells. To do so, we build on the 2SLS specification reported in column (5) of Table 5 and separately interact the relative AI automation and augmentation measures with baseline college share and occupational routine intensity.[18] We then plot the resulting marginal effects evaluated at selected percentiles of each characteristic's distribution. Panel A of Figure 6 shows modest positive estimated effects of

[18] We instrument these interaction terms using the corresponding interactions between the relative historical-breakthrough instruments and the respective characteristic. We also interact total automation and augmentation exposure with that characteristic, allowing the effects of overall exposure to vary along the same dimension.

Table 5: Effects of Relative AI Automation and Augmentation on Local Labor Markets

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
Relative AI Automation	-0.011	-0.006	-0.029***	-0.030***	0.005
	(0.014)	(0.012)	(0.008)	(0.008)	(0.009)
Relative AI Augmentation	0.003	0.025**	0.016**	0.009	0.005
	(0.008)	(0.010)	(0.008)	(0.007)	(0.008)
Panel B: Employment change					
Relative AI Automation	-0.000	0.000	-0.025***	-0.027***	-0.004
	(0.013)	(0.012)	(0.009)	(0.008)	(0.009)
Relative AI Augmentation	-0.022***	-0.010	-0.021***	-0.027***	-0.028***
	(0.007)	(0.009)	(0.007)	(0.007)	(0.008)
Panel C: Hourly wage change					
Relative AI Automation	-0.002	0.002	0.005*	0.005*	0.013***
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Relative AI Augmentation	0.017***	0.025***	0.032***	0.030***	0.028***
	(0.003)	(0.004)	(0.004)	(0.004)	(0.003)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	379.22	481.74	486.33	463.04	424.84
Observations	89,001	89,001	89,001	89,001	89,001

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using the common estimation sample of column (5). The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages between 2005–07 and 2022–24. Relative AI automation and augmentation exposure are defined as the AI share of total exposure within each channel and standardized within this sample. They are instrumented using the corresponding AI shares of distance-weighted historical-breakthrough exposure constructed from patents granted during 1980–2000 outside the two-ring exclusion buffer. The estimation sample requires positive total contemporary and historical-breakthrough exposure in both channels. All specifications control for total automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

relative AI automation exposure on employment, hourly-wage, and wage-bill growth at the lower end of the college-share distribution. These effects decline steadily as college share rises, with the employment and wage-bill effects turning negative toward the upper end of the distribution. This gradient suggests that shifting the composition of automating innovation from non-AI toward AI may be particularly labor-displacing for more highly educated workforces.

Panel B shows the opposite gradient across the college-share distribution, revealing substantial heterogeneity behind the offsetting average hourly-wage and employment effects reported in Table 5. Across the lower and middle parts of the college-share distribution, shifting augmentation exposure toward AI substantially dampens employment and wage-bill growth but produces

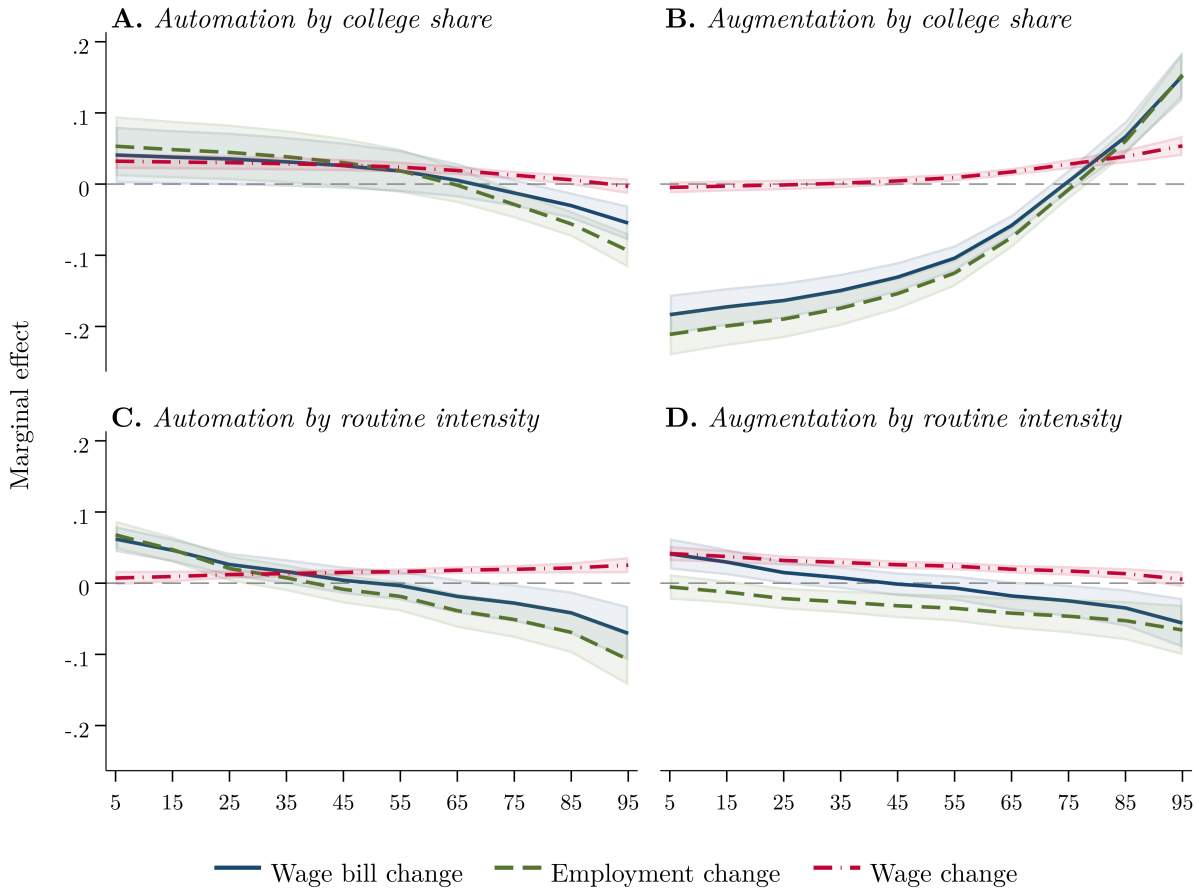


Figure 6: Marginal Effects of Relative AI Exposure across the College-Share and Routine-Intensity Distributions

Notes: The figure reports 2SLS marginal effects of a one-standard-deviation increase in relative AI automation and augmentation exposure, evaluated from the 5th through the 95th percentiles of baseline college share (Panels A and B) and occupational routine intensity (Panels C and D). The three series show effects on the log changes in the wage bill and FTE employment and the change in mean log hourly wage. Both relative AI exposure measures and their interactions with the respective characteristic are instrumented using the corresponding relative historical-breakthrough exposure measures and their interactions with that characteristic. All specifications control for baseline college share, occupational routine intensity, total automation and augmentation exposure, the interactions of both total-exposure measures with the respective characteristic, baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), and commuting-zone-by-broad-occupation fixed effects. Percentile cutoffs are calculated using 2005–07 FTE employment-share weights, which are also used in the regressions. The Kleibergen–Paap F-statistics are 55.6 for the college-share specification and 80.5 for the routine-intensity specification. Standard errors are clustered by commuting zone, and shaded bands denote 95 percent confidence intervals.

effectively no additional hourly-wage growth. As the college share rises, however, the estimated effects on employment and wage-bill growth eventually turn positive, while the wage effect becomes noticeably larger at the top of the distribution. This pattern is consistent with AI-enabled tools being more complementary to expertise and judgment in occupation–commuting-zone cells with more highly educated workforces. Taken together, Panels A and B show that both the adverse and favorable effects of shifting the innovation mix toward AI are unevenly distributed. Conse-

quently, the overall skill gradient of AI's labor-market effects depends on the relative strength of the automation and augmentation channels.

Panels C and D show the routine-intensity gradients of relative AI automation and augmentation exposure, respectively. The estimated employment and wage-bill effects of both channels become more negative as routine intensity rises, although this pattern is more pronounced for automation. The hourly-wage effect of relative AI automation moves in the opposite direction, becoming increasingly positive with routine intensity. For relative AI augmentation, all three effects decline as routine intensity rises, although the hourly-wage effect remains positive throughout. These results indicate that, in routine-intensive occupations, the positive wage effects of a higher AI share do not offset weaker employment growth. Overall, Figure 6 highlights that the labor-market consequences of AI are not only channel-specific but also vary strongly with workforce skill composition and occupational routine intensity.

7 Conclusion

This paper develops an empirical framework for studying how AI reshapes labor markets by distinguishing innovations that automate workers' tasks from those that augment their productive capabilities by improving the tools they use. Using transformer-based language models, we link text from geolocated U.S. patents to occupational task and tool descriptions and construct occupation–commuting-zone measures of exposure to both channels. We then exploit variation in contemporary local AI exposure predicted by the spatial diffusion of historical AI breakthroughs to estimate their effects on employment, hourly wages, and wage bills.

Consistent with task-based theory, the results reveal a clear contrast between the two channels. AI automation reduces employment and wage-bill growth, whereas AI augmentation increases wage bills, primarily through higher wages but more recently also through employment gains. These channel-specific effects become more pronounced over time and are robust across several alternative exposure measures and instrument constructions.

Our results further highlight that AI changes the strength and form of these two channels. As the AI share of automating innovation rises, the labor-displacing effects of automation become stronger. Within augmentation, by contrast, a higher AI share shifts adjustment away from employment growth and toward wage gains. These compositional effects are highly heterogeneous: when innovation shifts toward AI, more educated local workforces experience larger losses through automation but greater gains through augmentation, while routine-intensive occupations tend to fare worse, particularly through the automation channel.

Three limitations define the scope of our analysis. First, our measures capture exposure to locally assigned patented innovations rather than realized workplace adoption. Assignee locations provide a natural geographic anchor for where patented innovations originate, although the same technology may ultimately be implemented differently across workplaces. Our measures also necessarily focus on patented innovation and therefore do not capture other forms of AI diffusion, including imported, open-source, or off-the-shelf technologies. Second, our exposure measures cover innovations patented between 2005 and 2021. The analysis therefore does not fully capture the recent diffusion of generative AI, and the magnitude and distribution of these effects may evolve as AI technologies continue to develop and spread. Third, our design identifies local rather than aggregate national labor-market effects. Accordingly, whether these local employment and wage effects aggregate nationally or are offset by adjustment across occupations and regions remains an important question beyond the scope of our design.

Our study points to a clear conclusion: AI does not have a single labor-market effect. AI that automates workers' tasks reduces employment and wage-bill growth, whereas AI that improves the tools workers use raises hourly wages and can increase employment. Which force dominates will depend on the direction, diffusion, and adoption of future AI innovation. The automation–augmentation distinction may therefore become increasingly important as AI accounts for a larger share of technological change and appears to affect workers differently from non-AI innovation within both channels.

References

Abadie, Alberto, Susan Athey, Guido W. Imbens, and Jeffrey M. Wooldridge (2023). "When Should You Adjust Standard Errors for Clustering?" In: *The Quarterly Journal of Economics* 138.1, pp. 1–35. DOI: 10.1093/qje/qjac038.

Acemoglu, Daron (2021). *Harms of AI*. NBER Working Paper 29247. National Bureau of Economic Research. DOI: 10.3386/w29247.

— (2025). "The Simple Macroeconomics of AI". In: *Economic Policy* 40.121, pp. 13–58. DOI: 10.1093/epolic/eiae042.

Acemoglu, Daron and David H. Autor (2011). "Skills, Tasks and Technologies: Implications for Employment and Earnings". In: *Handbook of Labor Economics*. Ed. by Orley Ashenfelter and David Card. Vol. 4B. Elsevier, pp. 1043–1171. DOI: 10.1016/S0169-7218(11)02410-5.

Acemoglu, Daron, David H. Autor, Jonathon Hazell, and Pascual Restrepo (2022). "Artificial Intelligence and Jobs: Evidence from Online Vacancies". In: *Journal of Labor Economics* 40.S1, S293–S340. DOI: 10.1086/718327.

Acemoglu, Daron, David H. Autor, and Simon Johnson (2026). *Building Pro-Worker Artificial Intelligence*. NBER Working Paper 34854. National Bureau of Economic Research. DOI: 10.3386/w34854.

Acemoglu, Daron and Pascual Restrepo (2018). "The Race between Man and Machine: Implications of Technology for Growth, Factor Shares, and Employment". In: *American Economic Review* 108.6, pp. 1488–1542. DOI: 10.1257/aer.20160696.

— (2019). "Automation and New Tasks: How Technology Displaces and Reinstates Labor". In: *Journal of Economic Perspectives* 33.2, pp. 3–30. DOI: 10.1257/jep.33.2.3.

— (2020). "The Wrong Kind of AI? Artificial Intelligence and the Future of Labour Demand". In: *Cambridge Journal of Regions, Economy and Society* 13.1, pp. 25–35. DOI: 10.1093/cjres/rsz022.

Agrawal, Ajay K., John McHale, and Alexander Oettl (2026). *Enhancing Worker Productivity without Automating Tasks: A Different Approach to AI and the Task-Based Model*. NBER Working Paper 34781. National Bureau of Economic Research. DOI: 10.3386/w34781.

Autor, David H., Caroline Chin, Anna Salomons, and Bryan Seegmiller (2024). "New Frontiers: The Origins and Content of New Work, 1940–2018". In: *The Quarterly Journal of Economics* 139.3, pp. 1399–1465. DOI: 10.1093/qje/qjae008.

Autor, David H. and David Dorn (2013). "The Growth of Low-Skill Service Jobs and the Polarization of the U.S. Labor Market". In: *American Economic Review* 103.5, pp. 1553–1597. DOI: 10.1257/aer.103.5.1553.

Autor, David H., Frank Levy, and Richard J. Murnane (2003). "The Skill Content of Recent Technological Change: An Empirical Exploration". In: *The Quarterly Journal of Economics* 118.4, pp. 1279–1333. DOI: 10.1162/003355303322552801.

Autor, David H. and Neil Thompson (2025). "Expertise". In: *Journal of the European Economic Association* 23.4, pp. 1203–1271. DOI: 10.1093/jeea/jvaf023.

Beckhusen, Julia B. (2020). *Recent Changes in the Census Industry and Occupation Classification Systems*. American Community Survey Technical Paper ACS-TP 78. U.S. Census Bureau. URL: https://www.census.gov/library/publications/2020/demo/acs-tp78.html.

Bekamiri, Hamid, Daniel S. Hain, and Roman Jurowetzki (2024). "PatentSBERTa: A Deep NLP-Based Hybrid Model for Patent Distance and Classification Using Augmented SBERT". In: *Technological Forecasting and Social Change* 206, p. 123536. DOI: 10.1016/j.techfore.2024.123536.

Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen (Aug. 2026). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence*. Working Paper. Stanford Digital Economy Lab. URL: https://digitaleconomy.stanford.edu/publications/canaries-in-the-coal-mine/.

Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond (2025). "Generative AI at Work". In: *The Quarterly Journal of Economics* 140.2, pp. 889–942. DOI: 10.1093/qje/qjae044.

Brynjolfsson, Erik and Andrew McAfee (2015). "Will Humans Go the Way of Horses? Labor in the Second Machine Age". In: *Foreign Affairs* 94.4, pp. 8–14. URL: https://www.foreignaffairs.com/world/will-humans-go-way-horses.

Brynjolfsson, Erik, Tom Mitchell, and Daniel Rock (2018). "What Can Machines Learn, and What Does It Mean for Occupations and the Economy?" In: *AEA Papers and Proceedings* 108, pp. 43–47. DOI: 10.1257/pandp.20181019.

Chen, Mark A. and Joanna Wang (2024). *Displacement or Augmentation? The Effects of AI Innovation on Workforce Dynamics and Firm Value*. Working Paper. URL: https://www.aeaweb.org/conference/2025/program/paper/BBsK4Zkd.

Conley, Timothy G. (1999). "GMM Estimation with Cross-Sectional Dependence". In: *Journal of Econometrics* 92.1, pp. 1–45. DOI: 10.1016/S0304-4076(98)00084-0.

Cui, Kevin Zheyuan, Mert Demirer, Sonia Jaffe, Leon Musolff, Sida Peng, and Tobias Salz (2026). "The Effects of Generative AI on High-Skilled Work: Evidence from Three Field Experiments with Software Developers". In: *Management Science*. Articles in Advance. DOI: 10.1287/mnsc.2025.00535.

Decorte, Jens-Joris, Jeroen Van Hautte, Thomas Demeester, and Chris Develder (2021). "JobBERT: Understanding Job Titles through Skills". In: *arXiv preprint arXiv:2109.09605*. DOI: 10.48550/arXiv.2109.09605.

Decorte, Jens-Joris, Jeroen Van Hautte, Chris Develder, and Thomas Demeester (2025). "Efficient Text Encoders for Labor Market Analysis". In: *IEEE Access* 13, pp. 133596–133608. DOI: 10.1109/ACCESS.2025.3589147.

Dorn, David (2009). "Essays on Inequality, Spatial Interaction, and the Demand for Skills". PhD thesis. University of St. Gallen.

Felten, Edward W., Manav Raj, and Robert Seamans (2021). "Occupational, Industry, and Geographic Exposure to Artificial Intelligence: A Novel Dataset and Its Potential Uses". In: *Strategic Management Journal* 42.12, pp. 2195–2217. DOI: 10.1002/smj.3286.

Firpo, Sergio, Nicole M. Fortin, and Thomas Lemieux (2011). *Occupational Tasks and Changes in the Wage Structure*. IZA Discussion Paper 5542. Institute for the Study of Labor (IZA). DOI: 10.2139/ssrn.1778886.

Gathmann, Christina, Felix Grimm, and Erwin Winkler (2024). *AI, Task Changes in Jobs, and Worker Reallocation*. CESifo Working Paper 11585. CESifo. URL: https://www.ifo.de/en/cesifo/publications/2024/working-paper/ai-task-changes-jobs-and-worker-reallocation.

Hampole, Menaka, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller (2025). *Artificial Intelligence and the Labor Market*. NBER Working Paper 33509. Revised September 2025. National Bureau of Economic Research. DOI: 10.3386/w33509.

Huang, Yueling (2024). *The Labor Market Impact of Artificial Intelligence: Evidence from US Regions*. IMF Working Paper 2024/199. International Monetary Fund. DOI: 10.5089/9798400288555.001.

Humlum, Anders and Emilie Vestergaard (2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. NBER Working Paper 33777. Revised March 2026. National Bureau of Economic Research. DOI: 10.3386/w33777.

Jaffe, Adam B., Manuel Trajtenberg, and Rebecca Henderson (1993). "Geographic Localization of Knowledge Spillovers as Evidenced by Patent Citations". In: *The Quarterly Journal of Economics* 108.3, pp. 577–598. DOI: 10.2307/2118401.

Kalyani, Aakash, Nicholas Bloom, Marcela Carvalho, Tarek Hassan, Josh Lerner, and Ahmed Tahoun (2025). "The Diffusion of New Technologies". In: *The Quarterly Journal of Economics* 140.2, pp. 1299–1365. DOI: 10.1093/qje/qjaf002.

Kelly, Bryan, Dimitris Papanikolaou, Amit Seru, and Matt Taddy (2021). "Measuring Technological Innovation over the Long Run". In: *American Economic Review: Insights* 3.3, pp. 303–320. DOI: 10.1257/aeri.20190499.

Kogan, Leonid, Dimitris Papanikolaou, Lawrence D. W. Schmidt, and Bryan Seegmiller (2023). *Technology and Labor Displacement: Evidence from Linking Patents with Worker-Level Data*. NBER Working Paper 31846. National Bureau of Economic Research. DOI: 10.3386/w31846.

Marguerit, David (2025). "Augmenting or Automating Labor? The Effect of AI Development on New Work, Employment, and Wages". In: *arXiv preprint arXiv:2503.19159*. DOI: 10.48550/arXiv.2503.19159.

Pairolero, Nicholas A., Alexander V. Giczy, Gerard Torres, Tisa Islam Erana, Mark A. Finlayson, and Andrew A. Toole (2025). "The Artificial Intelligence Patent Dataset (AIPD) 2023 Update". In: *The Journal of Technology Transfer* 50, pp. 2587–2610. DOI: 10.1007/s10961-025-10189-8.

Prytkova, Ekaterina and Fabien Petit (2026). "The Employment Impact of Emerging Digital Technologies: Evidence from US Labor Markets". In: *The Changing Nature of Work*. Ed. by Susan N. Houseman, Anne Polivka, and Ayşegül Şahin. University of Chicago Press. Chap. 12. URL: https://www.nber.org/books-and-chapters/changing-nature-work/employment-impact-emerging-digital-technologies-evidence-us-labor-markets.

Roncone, Francesco (2026). *Within-Occupation Technological Change: Effects on Spatial Skill Sorting and Wage Inequality*. SSRN Working Paper 6986776. Social Science Research Network. DOI: 10.2139/ssrn.6986776.

Ruggles, Steven, Sarah Flood, Matthew Sobek, Daniel Backman, Grace Cooper, Julia A. Rivera Drew, Stephanie Richards, Renae Rodgers, Jonathan Schroeder, and Kari C. W. Williams (2025). *IPUMS USA: Version 16.0 [dataset]*. Minneapolis, MN: IPUMS. DOI: 10.18128/D010.V16.0.

Tolbert, Charles M. and Molly Sizer (1996). *U.S. Commuting Zones and Labor Market Areas: A 1990 Update*. ERS Staff Paper AGES-9614. U.S. Department of Agriculture, Economic Research Service. DOI: 10.22004/ag.econ.278812.

Webb, Michael (2019). *The Impact of Artificial Intelligence on the Labor Market*. SSRN Working Paper 3482150. Social Science Research Network. DOI: 10.2139/ssrn.3482150.

A Data Appendix

A.1 ACS Microdata

To analyze wage and employment changes at the occupation–commuting-zone level, sufficiently large cell sizes are required. We therefore pool American Community Survey (ACS) microdata from 2005–2007, 2014–2016, and 2022–2024, obtained from IPUMS (Ruggles et al. 2025). The underlying sample consists of employed individuals aged 18–64 years, excluding individuals residing in institutional group quarters, such as prisons and mental institutions, workers in military occupations, and unpaid family workers.

Commuting Zone Assignment. We assign each individual to one of the 741 commuting zones in the United States using the commuting zone classification developed by Tolbert and Sizer (1996). Because Census and ACS microdata identify individuals only at the Public Use Microdata Area (PUMA) level, we combine PUMAs with the widely used PUMA-to-CZ crosswalk provided by Dorn (2009). When a PUMA overlaps multiple commuting zones, we replicate observations across the relevant commuting zones and scale their sampling weights according to the share of the PUMA population residing in each commuting zone. The resulting weights sum to the original ACS sampling weight.

Weight Adjustments. We divide the ACS person sampling weights by three to account for pooling three survey years within each period. We then construct labor-supply weights by multiplying the adjusted person weights by annual hours worked, calculated as the product of weeks worked during the previous year and usual weekly hours worked. Individuals reporting zero weeks worked are assigned missing labor-supply weights and are therefore excluded from calculations using these weights.

Occupation–Commuting-Zone Cells. To ensure reliable estimates of employment and wage changes, we retain only occupation–commuting-zone cells containing at least ten observations in both the 2005–2007 and 2022–2024 periods used in the long-run analysis. This restriction removes 199,283 of the 296,938 occupation–commuting-zone cells (67.1 percent), but only 2.44 million of the 21.04 million individuals in the underlying sample (11.6 percent). Despite the exclusion of a noticeable number of small cells, all 430 occupations in the *occ2010fr* occupation panel remain represented in the final long-run estimation sample. In unreported robustness checks, increasing the minimum cell size to 20 or 30 leaves the main findings intact.

Employment, Wages, and Wage Bills. Our main outcome variables are changes in employment, hourly wages, and wage bills at the occupation–commuting-zone level between 2005–07 and 2022–24. Employment is measured in full-time-equivalent (FTE) workers to account for changes in both employment and hours. For each worker, annual hours are calculated as weeks worked multiplied by usual weekly hours. Following Autor et al. (2024), we construct occupation–

commuting-zone FTE employment by summing individual annual hours using the adjusted ACS person sampling weights and dividing the resulting total by 35 hours per week and 50 weeks per year. Occupation–commuting-zone wage bills are constructed by aggregating workers' annual labor earnings, measured in 2010 CPI-adjusted dollars, using the adjusted person weights. Individual hourly wages are calculated by dividing annual labor earnings by annual hours worked and are winsorized at the 1st and 99th percentiles separately for each period to limit the influence of outliers. Average log hourly wages within occupation–commuting-zone cells are constructed using the labor-supply weights described above. The empirical analysis is based on log changes in employment and wage bills and changes in average log hourly wages.

A.2 Occupation Panel Construction

The analysis uses the balanced occupation panel, *occ2010fr*, introduced in Roncone (2026), which consists of 430 occupations that can be consistently tracked across Census and ACS occupation classifications from 2000 onward. The panel was constructed to preserve occupational granularity while maintaining a stable occupation structure over time. Relative to the widely used balanced occupation panel developed by Dorn (2009), the *occ2010fr* panel preserves approximately 30 percent more occupation categories. The additional detail is particularly valuable for technology- and knowledge-intensive occupations, which are represented by a larger and more differentiated set of occupation categories in post-2000 classifications.

The panel was constructed by combining Census and IPUMS occupation crosswalks with official occupation conversion rates (Beckhusen 2020). For occupation splits and newly emerging occupations, mappings follow the largest employment-weighted conversion rate when the corresponding share exceeds 80 percent; otherwise, a broader occupation category is constructed to preserve consistency across classification systems. Occupation mergers are mapped using the official conversion rates. Residual occupations classified as "not elsewhere classified" (n.e.c.) are grouped within broader 2-digit or 3-digit occupation categories when a consistent assignment to more detailed occupations is not possible across classification systems.

A.3 Occupational Routine Intensity

Following Firpo et al. (2011), we construct an occupation-level routine-intensity index using five O*NET version 16.0 Work Context measures: degree of automation, importance of repeating the same tasks, structured versus unstructured work, pace determined by the speed of equipment, and time spent making repetitive motions. We map these measures to the *occ2010fr* panel using the same occupational crosswalk used to assign O*NET tasks and tools. Each component is min–max normalized to range from 0 to 10 across the 430 panel occupations, with the structured-versus-unstructured-work measure reversed so that higher values indicate more structured work. We take the unweighted mean of the five components and rescale the resulting index to range from 0 to 10, with higher values indicating greater routine intensity. The resulting occupation-level value is assigned to every corresponding occupation–commuting-zone cell.

A.4 Assignment of O*NET Tasks and Tools

Occupational tasks are obtained from the O*NET *Task Ratings* database, while occupational tools are obtained from the O*NET *Tools Used* database. We use O*NET version 25.0 as the fixed reference inventory. The *Tools Used* inventory covers a broad range of productive equipment, machinery, instruments, and computing devices across occupations, including desktop computers, drainage pumps, hammers, and viscometers. O*NET also provides a separate *Technology Skills* inventory, which primarily records named software packages, programming languages, and digital applications associated with each occupation. We exclude this inventory because our framework defines augmentation through the tools workers use rather than through their technology-related skill requirements. Including it could blur this distinction and cause the augmentation measure to partly capture pre-existing differences in occupational digital specialization.

Dealing with Newly Introduced O*NET Occupations. Some detailed occupations in O*NET version 25.0 were introduced only after the beginning of our AI exposure period. These additions may reflect either genuinely emerging occupations or changes in the taxonomy that recognize occupations and specializations that previously existed but were not separately classified. Because these cases cannot be consistently distinguished, including all later additions could make measured exposure partly dependent on subsequent changes in the O*NET taxonomy. To reduce this influence, we retain version 25.0 task and tool records only for occupations that can be linked to occupations already present in version 16.0, approximately a decade earlier. The resulting task and tool inventories draw on more than 850 detailed O*NET occupations.

Aggregating Multiple O*NET Occupations. Although the 8-digit O*NET occupation codes are more granular than the occupation categories used in the ACS, the two systems can be linked through the Standard Occupational Classification (SOC) framework. To translate the O*NET task and tool inventories to our *occ2010fr* occupation panel, we aggregate information across matched O*NET occupations. Whenever O*NET decomposes a 6-digit occupation into multiple specialized 8-digit categories, we combine the task and tool descriptions from these specialized categories with those of the corresponding broader 6-digit occupation to retain as much information as possible.

Task and Tool Weighting. To account for differences in the relative importance of tasks within occupations, we use the task importance ratings provided in O*NET version 25.0. These ratings are measured on a five-point scale ranging from 1 (not important) to 5 (extremely important). Although O*NET tasks are generally occupation-specific, duplicate tasks can arise when closely related O*NET occupations are aggregated into the same occupation in our panel. In these cases, we collapse duplicate task descriptions into a single task and average the contributing importance ratings. We then normalize the resulting task importance ratings within each panel occupation so that the weights sum to one, thereby accounting for differences across occupations in the number of recorded tasks.

As O*NET does not provide analogous importance ratings for tools, we construct occupation-specific tool weights by measuring the semantic similarity between each O*NET tool description and the unit-normalized, importance-weighted task profile of the corresponding occupation, constructed using these task weights. We calculate these similarities using JobBERT V2, a language model trained specifically on job- and skill-related text, set negative similarities to zero, and normalize the resulting scores across tools within each occupation so that the weights sum to one. In an alternative construction, we replace the task profile with a profile based on each occupation's O*NET Alternate Titles and measure the semantic similarity between this title-based representation and its tools. This alternative addresses the concern that using task information to weight tools may mechanically alter the relationship between our task-automation and tool-augmentation measures. Although both weighting schemes yield nearly identical estimates, we prefer the task-profile weighting scheme because its economic rationale is more direct.

Dealing with Residual Occupation Categories. We distinguish between two types of residual occupations in our panel. First, some residual occupations ("other") can be directly matched to residual O*NET occupations and therefore inherit the corresponding task and tool descriptions. Second, some residual occupation categories ("n.e.c.") either do not correspond to any O*NET occupation retained in our inventory or can be matched only to a single 8-digit O*NET occupation that is unlikely to be representative of the broader residual category. In these cases, we do not assign tasks and tools directly. Instead, after constructing the automation and augmentation exposure measures for all 416 occupations with valid O*NET matches, we impute the corresponding measures for the 14 remaining residual occupations within each commuting zone using the unweighted mean across directly matched occupations in the corresponding broader occupational category.

Resulting Task and Tool Inventories. After mapping O*NET version 25.0 to the *occ2010fr* occupation panel, the resulting inventories for occupations with valid O*NET matches contain approximately 16,100 distinct task descriptions and 17,300 distinct tool examples. Directly matched occupations are assigned, on average, 39 tasks (range: 6–264) and 80 tools (range: 3–676). In contrast to tasks, individual tools are often shared across a much broader range of occupations. For example, whereas *jet guns* are assigned only to pumping station operators, *desktop computers* appear in more than 250 occupations in our panel. The semantic weighting procedure described above allows the relevance of the same tool to vary across occupations by assigning occupation-specific weights. As a result, *desktop computers* receive the highest semantic weight among the tools assigned to financial analysts but rank 217th out of 232 tools for printing press operators. For illustration, Table A.1 reports the five highest-ranked tasks and tools for four selected occupations.

Table A.1: Examples of Top Tasks and Tools for Selected Detailed Occupations

Top 5 Tasks	*Top 5 Tools*
A. Lodging Managers (340)	
1. Answer inquiries on hotel policies and services, and resolve guest complaints.	1. Cash registers
2. Set room rates, establish budgets, and allocate departmental funds.	2. Credit card processing machines
3. Coordinate hotel activities with other managers.	3. Multi-line telephone systems
4. Greet and register guests.	4. 10-key calculators
5. Monitor hotel revenue.	5. Private automatic branch exchange (PABX) systems
B. Aerospace Engineers (1320)	
1. Develop mathematical models for engineering design.	1. Impinging jet apparatus
2. Conduct experimental and stress tests on aircraft and aerospace systems.	2. Spectrometers
3. Develop conceptual designs for aircraft and aerospace systems.	3. Axial flow turbines
4. Investigate and resolve technical problems with aircraft and aerospace vehicles.	4. Flight management systems (FMS)
5. Prepare technical reports and documentation for engineering staff and customers.	5. Propellant combustion chambers
C. Personal and Home Care Aides (4610)	
1. Provide bedside and personal care, including mobility and hygiene assistance.	1. Wheelchairs
2. Maintain client records and report changes in client condition.	2. Walkers
3. Monitor vital signs and administer healthcare tasks under clinical supervision.	3. Tablet computers
4. Participate in case reviews to evaluate client needs and plan ongoing care.	4. Back braces
5. Provide personal care and companionship during recovery or family disruption.	5. Shower chairs
D. Tire Builders (8940)	
1. Place tires into molds for new tread.	1. Tire stitchers
2. Fit inner tubes and final layers of rubber onto tires.	2. Curing presses
3. Buff tires according to specifications for width and undertread depth.	3. Tire band building machines
4. Build rubber treads onto tire casings for vulcanization.	4. Rubber slab cutters
5. Apply solvents to tire plies to ensure adhesion and strengthen tires.	5. Tire building drums

Notes: The table reports the five highest-ranked tasks and tools for selected occupations after mapping O*NET version 25.0 to the *occ2010fr* occupation panel. Tasks are ranked according to normalized O*NET importance weights. Tools are ranked according to the baseline task-profile weights described in the text. The corresponding weights are omitted for brevity. Task descriptions are abbreviated for presentation.

A.5 Patent Data

We construct the patent-level dataset from USPTO PatentsView records.[19] Starting from the full set of granted patents, we retain utility patents granted between 1980 and 2021, exclude reissue and withdrawn patents, and merge patent titles and abstracts using the patent identifier. We remove all patents without an abstract. This restriction affects only 796 patents and therefore represents a negligible share of the underlying patent universe. The resulting patent-level dataset contains unique patent identifiers and complete title and abstract information for all retained observations.

Next, we merge the cleaned patent records to the Artificial Intelligence Patent Dataset 2023 (AIPD 2023) using the patent identifier (Pairolero et al. 2025). Our baseline indicator classifies a patent as AI-related if its predicted probability for at least one of eight AI functions is at least 86 percent: machine learning, evolutionary computation, natural language processing, speech, computer vision, knowledge processing, AI hardware, and planning and decision systems. The component classifications are not mutually exclusive, so a patent may satisfy the 86-percent threshold for more than one AI component.

We additionally merge the patent records to the breakthrough measures of Kelly et al. (2021) using the same patent identifier. Our baseline breakthrough indicator is based on their residual forward-similarity measure, which identifies patents that are relatively dissimilar to preceding inventions but closely related to patents granted over the subsequent ten years. Following Autor et al. (2024), we classify patents in the top decile of the corresponding grant-year distribution as breakthroughs. In the empirical analysis, AI breakthrough patents granted between 1980 and 2000 are used to construct the historical instruments, whereas all eligible patents granted between 2005 and 2021 are used to construct contemporaneous local exposure. Patents granted between 2001 and 2004 therefore enter neither measure by construction, although they remain in the underlying patent dataset for descriptive exercises.

To assign patents geographically, we merge assignee and inventor location information and map U.S. county identifiers from the USPTO data to commuting zones. We use the location of the first listed assignee whenever it can be mapped to a U.S. commuting zone and, if no valid assignee commuting-zone assignment is available, the location of the first listed inventor. Among geographically assigned domestic patents with assignee information, 98.4 percent have a single assignee, while approximately 14 percent of geographically assigned domestic patents rely on the inventor-location fallback. Foreign patents and patents that cannot be assigned to a valid U.S. commuting zone under this assignee–inventor assignment rule are excluded when constructing local U.S. patent exposure.

The resulting geographically assigned domestic patent universe contains 48,452 AI patents in the 1980–2000 historical instrument window, of which 28,763 satisfy our breakthrough criterion, and 470,498 AI patents in the 2005–2021 contemporaneous exposure window. Table A.2 reports these AI patent counts across the eight AIPD components.

[19] PatentsView is an official U.S. Patent and Trademark Office research database providing structured and disambiguated patent data.

Table A.2: AI Patents and AI Breakthrough Patents by AI Component and Period

	Instrument window, 1980–2000				Exposure window, 2005–2021	
	AI patents		AI breakthroughs		AI patents	
AI component	*N*	Share (%)	*N*	Share (%)	*N*	Share (%)
Planning and decision systems	14,111	29.1	8,016	27.9	212,535	45.2
AI hardware	18,147	37.5	12,047	41.9	191,219	40.6
Knowledge processing	4,681	9.7	4,038	14.0	128,238	27.3
Computer vision	13,829	28.5	8,604	29.9	103,986	22.1
Natural language processing	5,327	11.0	4,137	14.4	83,451	17.7
Machine learning	4,580	9.5	3,078	10.7	65,139	13.8
Speech	4,591	9.5	2,474	8.6	33,923	7.2
Evolutionary computation	1,048	2.2	629	2.2	21,888	4.7
Any AI component	48,452	100.0	28,763	100.0	470,498	100.0

Notes: The table reports AI patents and AI breakthrough patents by component of the USPTO Artificial Intelligence Patent Dataset (Pairolero et al. 2025) for the historical instrument window (1980–2000) and the contemporaneous exposure window (2005–2021). A patent is assigned to a component when its predicted probability for that component is at least 86 percent. AI breakthrough patents are AI-classified patents in the top decile of the Kelly et al. (2021) breakthrough measure within the corresponding grant-year distribution. Component classifications are not mutually exclusive, so patents satisfying the threshold for multiple components are counted in each applicable row. Consequently, component-specific counts and shares may sum to more than the number of unique AI patents and 100 percent, respectively. Shares are calculated relative to the number of unique patents classified in at least one AI component in the corresponding column block; the "Any AI component" row therefore equals 100 percent by construction. Patent counts are restricted to domestic patents with a valid U.S. commuting-zone assignment.

B Semantic Content of AI Automation and Augmentation

The word clouds summarize the semantic content of AI patents most closely related to occupational tasks and tools. Within each AI function, patents granted between 2005 and 2021 are ranked separately by their maximum semantic similarity to any occupational task for automation and to any occupational tool for augmentation. The 5,000 highest-scoring patents in each function–channel combination are used to construct the clouds from patent titles and abstracts. Automation clouds display lemmatized verb–object pairs, while augmentation clouds display distinctive words and two-word phrases. Term size reflects how distinctive a term is for a given AI function relative to the other AI functions within the same channel, rather than how frequently the term appears in absolute terms.

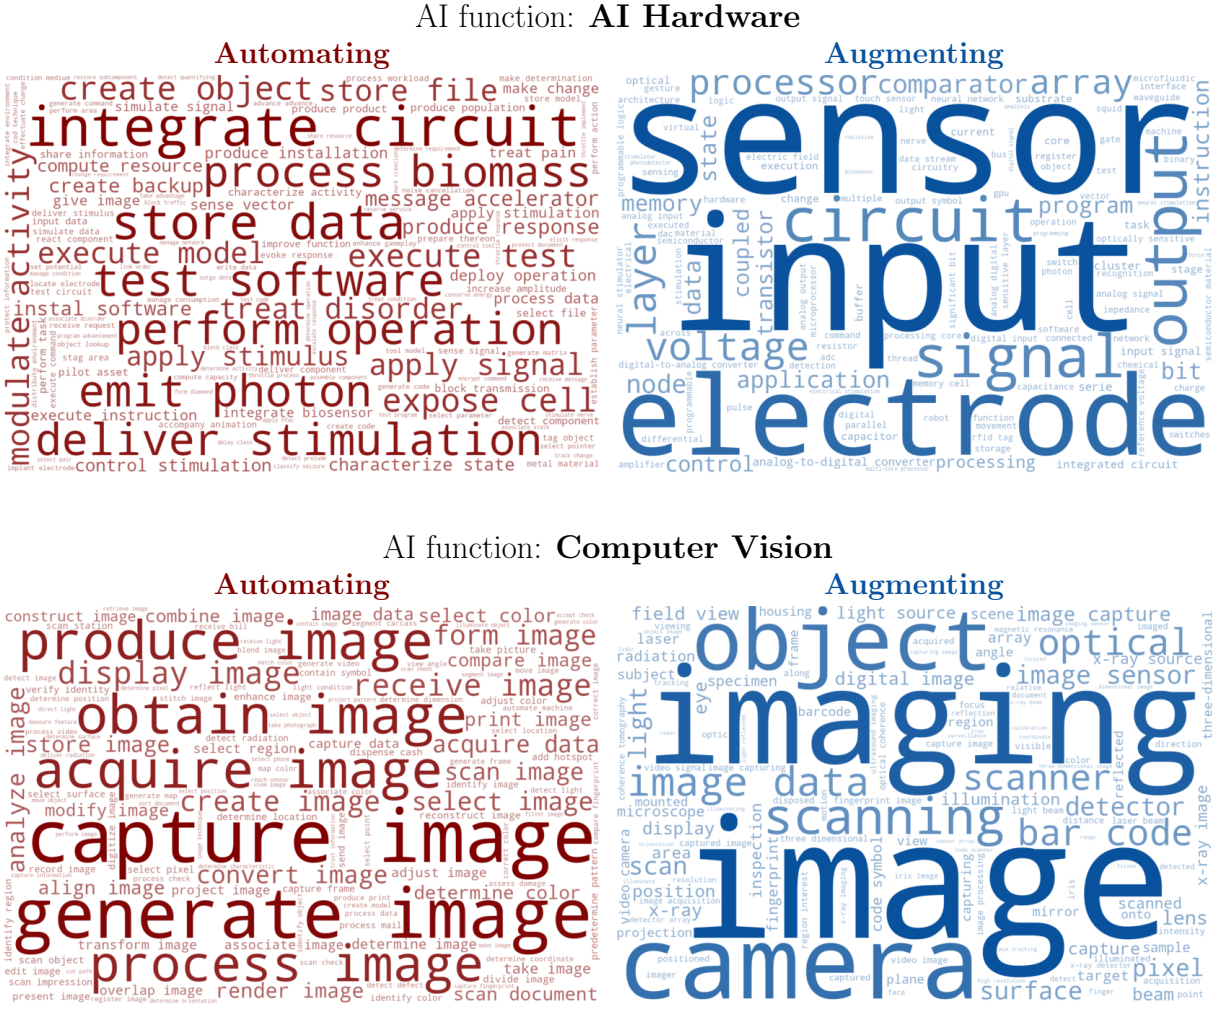
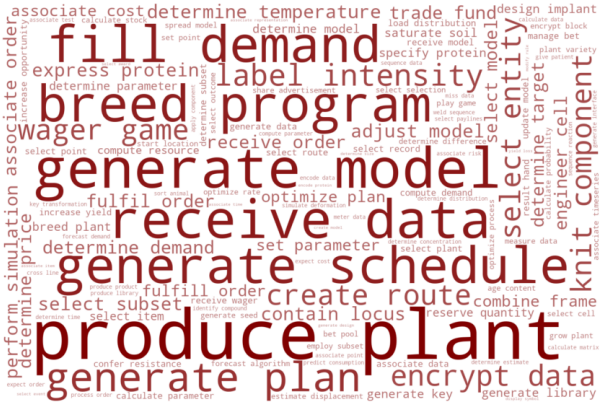


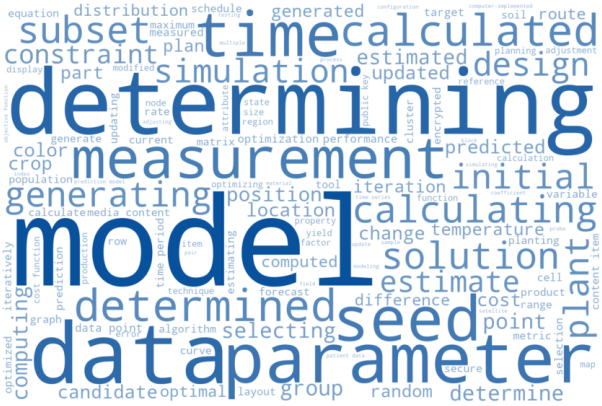
Figure B.1: Word Clouds of Automating and Augmenting AI Patents by AI Function

AI function: **Evolutionary Computation**

Automating



Augmenting

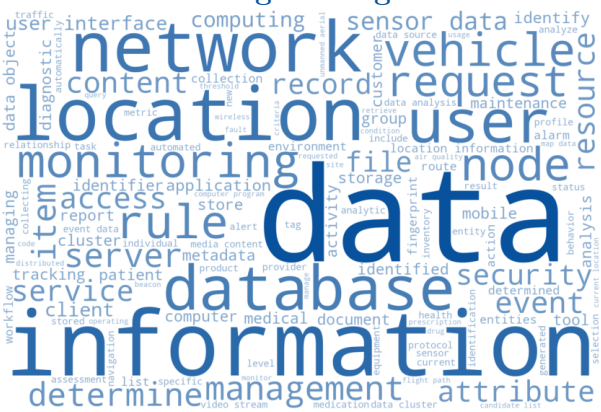


AI function: **Knowledge Processing**

Automating



Augmenting



AI function: **Machine Learning**

Automating



Augmenting

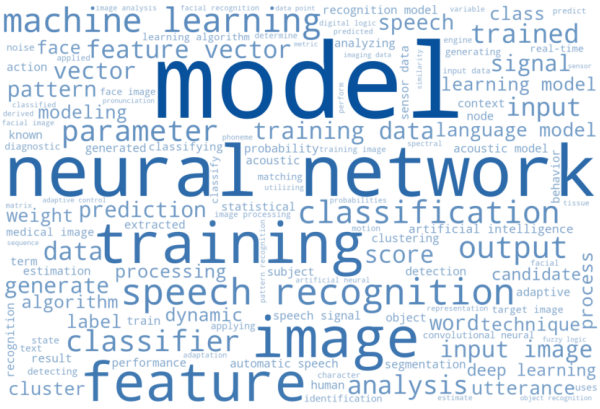


Figure B.1 continued.

AI function: **Natural Language Processing**

Automating



Augmenting

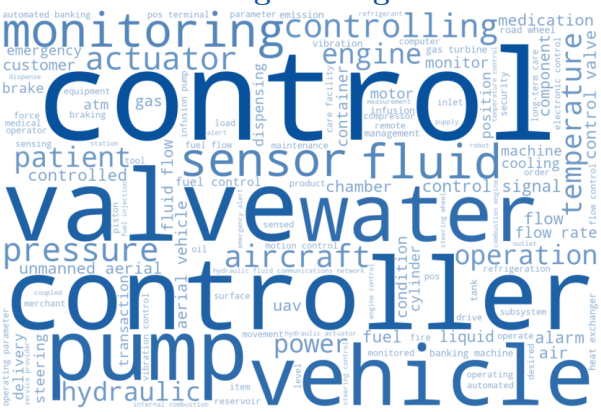


AI function: **Planning & Decision Systems**

Automating



Augmenting



AI function: **Speech**

Automating



Augmenting



Figure B.1 continued.

C Additional Descriptive Evidence

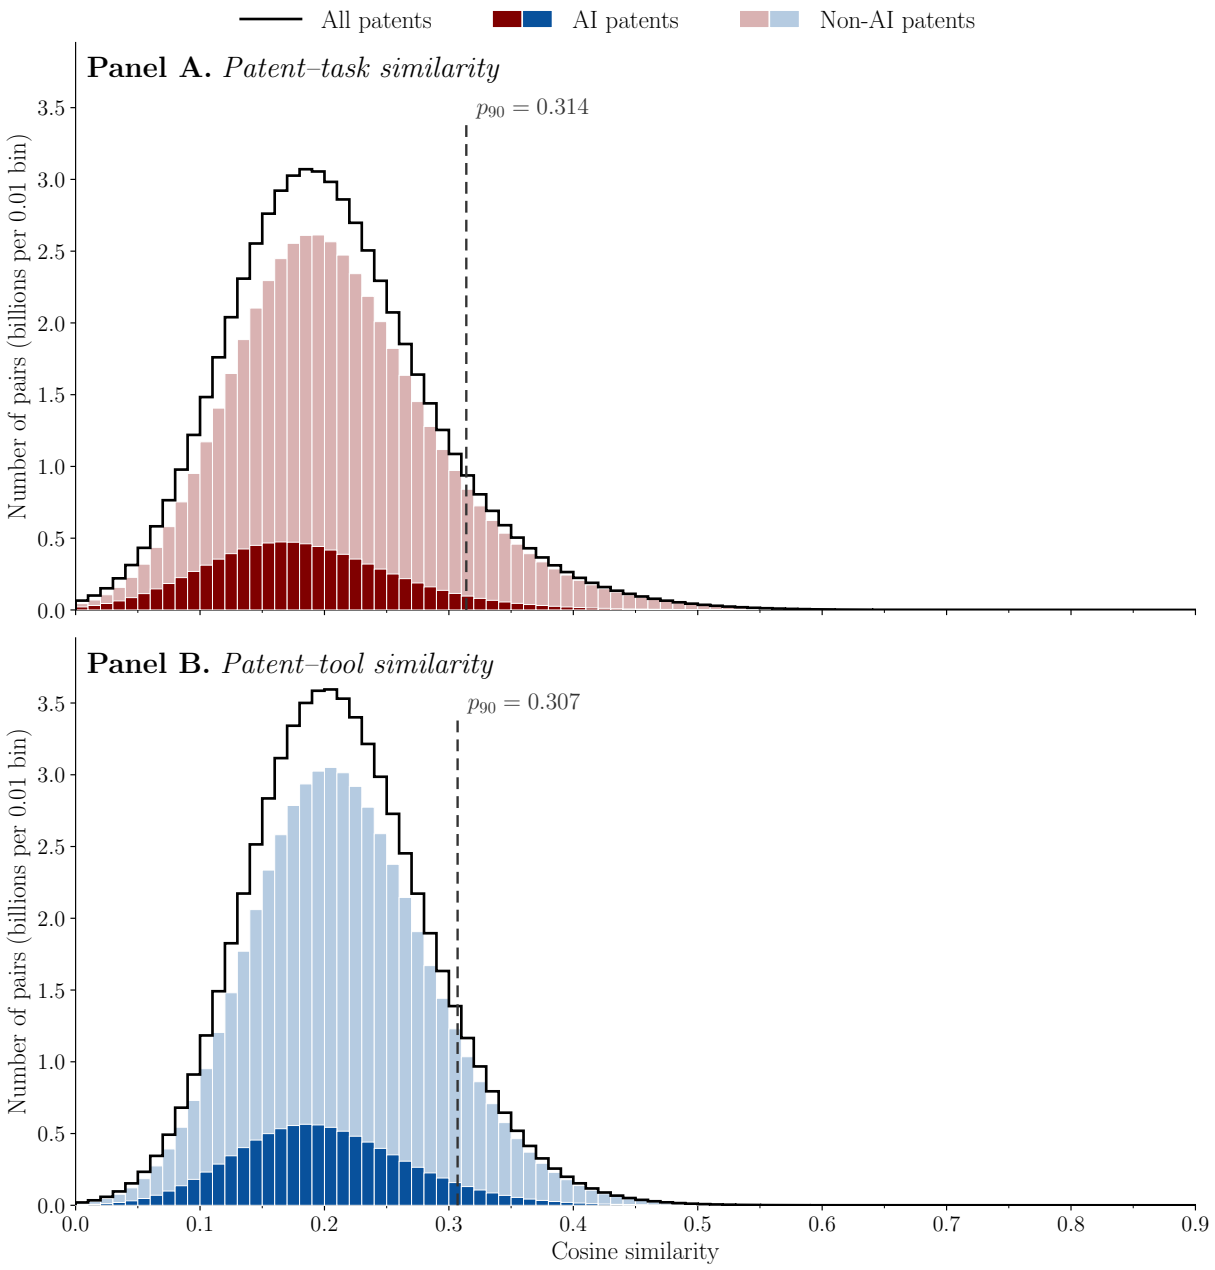


Figure C.1: Distributions of Patent–Task and Patent–Tool Cosine Similarity by AI Status

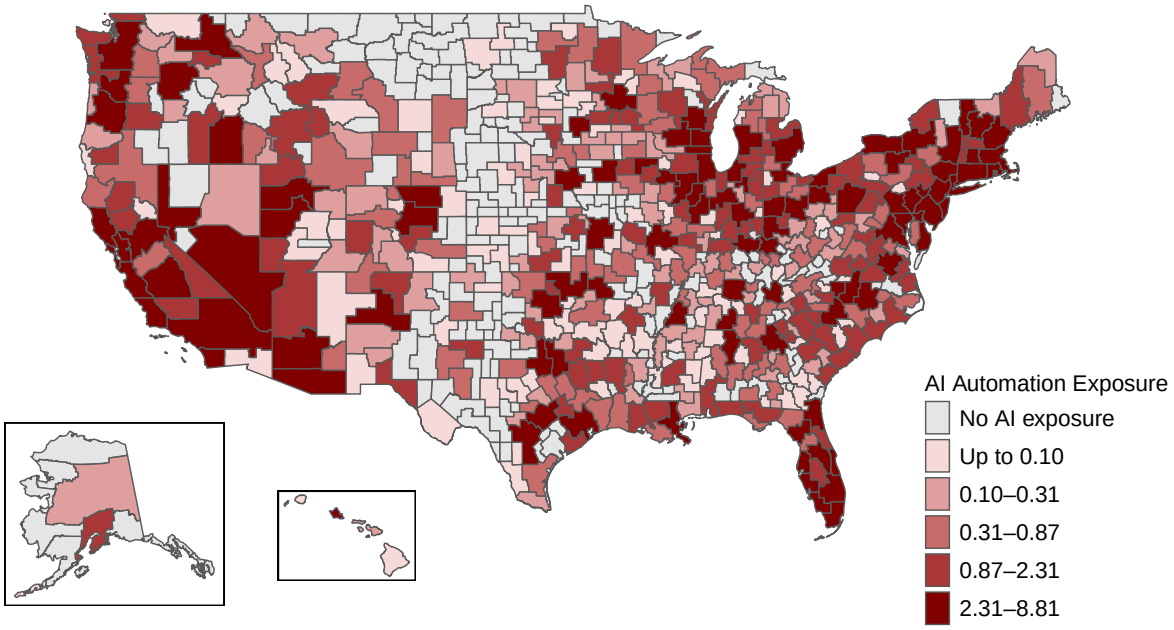
Notes: The figure plots cosine-similarity distributions for patent–task pairs (Panel A, JobBERT V2) and patent–tool pairs (Panel B, PatentSBERTa), separately by AI status. The sample comprises geographically assigned U.S. utility patents granted during 1980–2021, paired with all task or tool descriptions assigned to occupations in the *occ2010fr* panel. The vertical dashed lines mark the 90th percentiles of the corresponding pooled similarity distributions used as the baseline matching thresholds. Similarity scores are grouped into bins in the figure, whereas the matching thresholds are calculated from the full distributions. Black outlines show the pooled distributions for all patent pairs, while the shaded distributions distinguish AI and non-AI patents using the panel-specific automation and augmentation color schemes.

Table C.1: Occupations with the Highest AI Automation and Augmentation Exposure

	AI Automation		AI Augmentation	
Panel A: AI exposure				
Rank	**Occupation**	**Exposure**	**Occupation**	**Exposure**
1	Statistical assistants	8.278	Radiation therapists	8.771
2	Proofreaders and copy markers	8.245	Food and tobacco roasting, baking, and drying machine operators and tenders	8.675
3	Astronomers and physicists	8.195	Ambulance drivers and attendants, except emergency medical technicians	8.653
4	Television, video, and motion picture camera operators and editors	8.173	Agricultural and biomedical engineers	8.639
5	Etchers and engravers	8.079	Proofreaders and copy markers	8.612
6	Rail-track laying and maintenance equipment operators	8.052	Agents and business managers of artists, performers, and athletes	8.605
7	Mathematicians, statisticians and mathematical science occupations	7.961	Buyers and purchasing agents, farm products	8.477
8	New accounts clerks	7.947	Economists	8.442
9	Computer hardware engineers	7.919	Actors	8.393
10	Biological technicians	7.891	Producers and directors	8.287
Panel B: Relative AI exposure				
Rank	**Occupation**	**AI share**	**Occupation**	**AI share**
1	Financial examiners	47.1%	Actuaries	49.1%
2	Personal financial advisors	46.2%	Economists	48.4%
3	Credit analysts	46.0%	Financial examiners	46.8%
4	Computer and information systems managers	45.3%	Sales engineers	46.8%
5	Database administrators and architects	45.0%	Operations research analysts	46.5%
6	Budget analysts	42.9%	Financial analysts	44.7%
7	Financial specialists, other	42.2%	Insurance underwriters	44.6%
8	Actuaries	42.0%	Budget analysts	43.6%
9	Public relations and fundraising managers	39.5%	Purchasing managers	42.9%
10	Urban and regional planners	38.4%	Statistical assistants	41.4%

Notes: The table reports the ten detailed occupations with the highest employment-weighted average AI automation and augmentation exposure across commuting zones, using 2005–07 FTE employment shares as weights. Panel A ranks occupations using inverse-hyperbolic-sine-transformed but unstandardized AI exposure. Panel B ranks occupations using relative AI exposure, defined as the AI share of total exposure within the corresponding channel. Relative exposure is averaged across occupation–commuting-zone cells with positive total exposure in the relevant channel. Rankings exclude residual n.e.c. occupations whose exposure measures are imputed from related occupations.

Panel A. *AI automation exposure*



Panel B. *AI augmentation exposure*

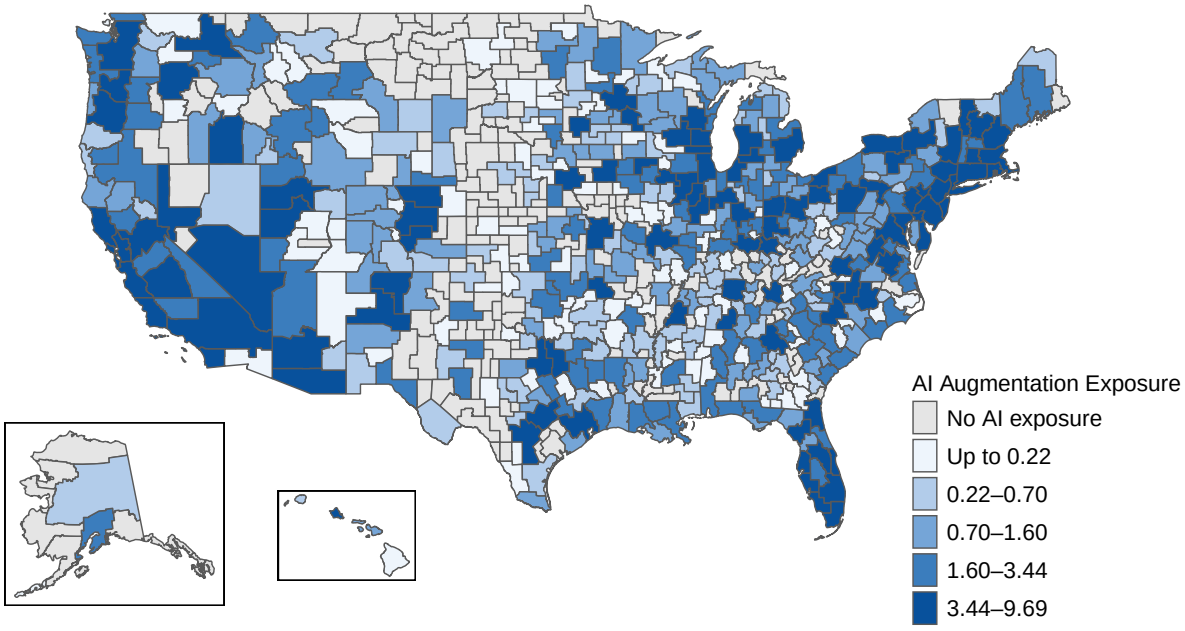
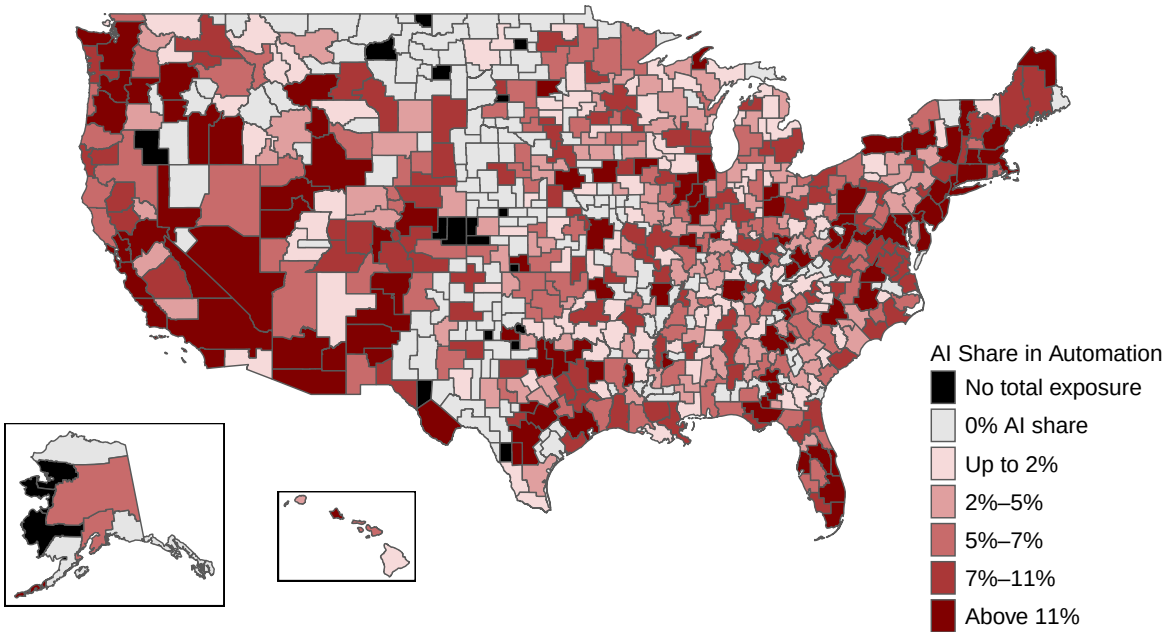


Figure C.2: Spatial Distribution of AI Automation and Augmentation Exposure

Notes: The maps display commuting-zone averages of the inverse-hyperbolic-sine-transformed AI automation and augmentation exposure measures, weighted across occupations using 2005–07 FTE employment shares. Commuting zones with zero exposure are shown in gray. Among commuting zones with strictly positive exposure, values are grouped into five quintiles calculated without employment weights.

Panel A. *Relative AI automation exposure*



Panel B. *Relative AI augmentation exposure*

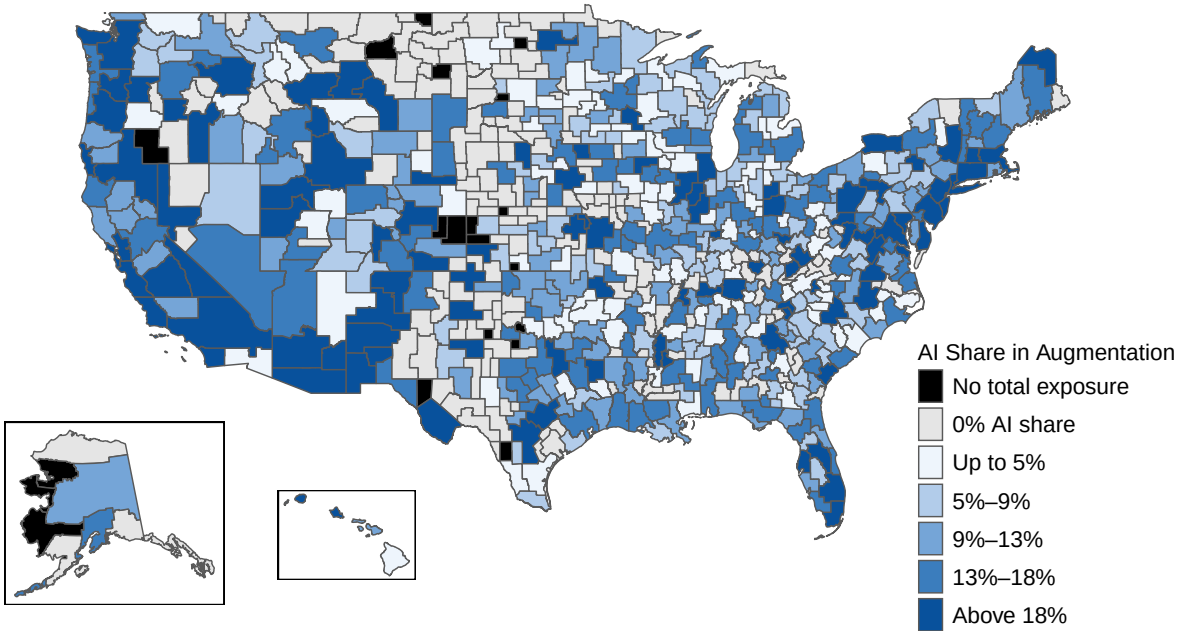


Figure C.3: Spatial Distribution of Relative AI Automation and Augmentation Exposure

Notes: The maps display commuting-zone averages of the relative AI automation and augmentation exposure measures, weighted across occupations using 2005–07 FTE employment shares. Commuting zones with a zero AI share are shown in gray. Among commuting zones with strictly positive relative exposure, values are grouped into five quintiles calculated without employment weights. Commuting zones for which the relative AI share is undefined because no total exposure is observed in the corresponding channel are shown in black.

D Complementary Evidence on Automation and Augmentation in the U.S. Patent Stock

D.1 Patent-Level Classification

This appendix develops a patent-level classification that assigns each U.S. utility patent granted between 1980 and 2021 to one of four mutually exclusive categories: *Automating*, *Augmenting*, *Dual*, or *Neither*. The classification is based on patents' semantic similarity to occupational tasks and tools and complements the occupation–commuting-zone exposure measures developed in Section 3. Whereas those measures aggregate many patent–task and patent–tool matches, the classification developed here serves a purely descriptive purpose by assigning each patent to only a single category based on its strongest task and tool relationships. Its implementation therefore differs from the pairwise thresholding procedure used to construct the exposure measures. The resulting classification allows us to characterize the evolution of automating and augmenting innovation in the U.S. patent stock and the changing role of AI within each channel.

D.1.1 Embedding and Similarity

The classification uses the same language models employed in the main analysis. For each patent p, we combine its title and abstract into a single text and embed it once using each model. Each occupational task $t = 1, \ldots, T$ is embedded using JobBERT V2 (Decorte et al. 2025), and each occupational tool $r = 1, \ldots, R$ is embedded using PatentSBERTa (Bekamiri et al. 2024). Let $\mathbf{e}_t^J$ and $\mathbf{e}_p^J$ denote the embeddings of task t and patent p, and let $\mathbf{e}_r^P$ and $\mathbf{e}_p^P$ denote the embeddings of tool r and patent p. The corresponding cosine similarities are

$$\begin{aligned} s_{tp} &= \cos\left(\mathbf{e}_t^J, \mathbf{e}_p^J\right), \quad t = 1, \ldots, T, \\ s_{rp} &= \cos\left(\mathbf{e}_r^P, \mathbf{e}_p^P\right), \quad r = 1, \ldots, R. \end{aligned} \tag{24}$$

D.1.2 Per-Patent Maximum Similarity and Threshold Selection

For this descriptive exercise, we assign each patent to a single, mutually exclusive category. We therefore summarize each patent's relationship to occupational tasks and tools using its maximum similarity to any task and any tool:

$$s_{p,\max}^{\text{task}} = \max_t s_{tp}, \qquad s_{p,\max}^{\text{tool}} = \max_r s_{rp}. \tag{25}$$

We define the patent-level classification thresholds $\tau_{\max}^{\text{task}}$ and $\tau_{\max}^{\text{tool}}$ as the 90th percentiles of the empirical distributions of $s_{p,\max}^{\text{task}}$ and $s_{p,\max}^{\text{tool}}$, respectively.

The max subscript distinguishes these thresholds, which are defined over per-patent maximum similarities, from the pairwise thresholds τ^{task} and τ^{tool} used to construct the occupation-level exposure measures. Applying the pairwise thresholding procedure used for the exposure measures would be inappropriate in this context. Because each patent is compared with thousands of occupational tasks and tools, almost all patents would exceed the task threshold for at least one

task and simultaneously the tool threshold for at least one tool, causing the majority of patents to be classified as *Dual*. Applying thresholds to the distributions of per-patent maximum similarities instead produces a more informative patent-level classification for the descriptive analysis of the patent stock that follows.

D.1.3 Classification Rule

Each patent is assigned to exactly one category according to whether its maximum task and tool similarities exceed the corresponding patent-level thresholds:

$$\text{Class}(p) = \begin{cases} \textit{Dual}, & \text{if } s_{p,\max}^{\text{task}} \geq \tau_{\max}^{\text{task}} \text{ and } s_{p,\max}^{\text{tool}} \geq \tau_{\max}^{\text{tool}}, \\ \textit{Automating}, & \text{if } s_{p,\max}^{\text{task}} \geq \tau_{\max}^{\text{task}} \text{ and } s_{p,\max}^{\text{tool}} < \tau_{\max}^{\text{tool}}, \\ \textit{Augmenting}, & \text{if } s_{p,\max}^{\text{task}} < \tau_{\max}^{\text{task}} \text{ and } s_{p,\max}^{\text{tool}} \geq \tau_{\max}^{\text{tool}}, \\ \textit{Neither}, & \text{if } s_{p,\max}^{\text{task}} < \tau_{\max}^{\text{task}} \text{ and } s_{p,\max}^{\text{tool}} < \tau_{\max}^{\text{tool}}. \end{cases} \tag{26}$$

By construction, approximately 10 percent of patents exceed the patent-level task threshold and 10 percent exceed the patent-level tool threshold. The share classified as *Dual* depends on the overlap between the upper tails of the two maximum-similarity distributions.

D.2 Automation and Augmentation in the U.S. Patent Stock

The patent-level classification reveals three descriptive patterns. First, AI becomes increasingly important within both automating and augmenting innovation. The lower panel of Figure D.1 shows that the AI shares within both classes are small during the 1980s but rise sharply from the mid-1990s onward. By 2021, AI accounts for approximately 14 percent of automating patents and 21 percent of augmenting patents. The AI share is higher within the augmenting class for most of the period, and the gap has widened in recent years. This pattern parallels the occupation-level evidence that AI accounts for a larger share of augmentation exposure than of automation exposure across most occupations.

Second, the upper panel of Figure D.1 shows that the annual shares classified exclusively as automating or augmenting decline over time. We interpret this pattern with caution because the patent-level thresholds are defined from the pooled 1980–2021 distributions of per-patent maximum task and tool similarities. Annual classification rates can therefore shift as patent descriptions become more or less semantically similar to the fixed O*NET task and tool inventories and as the overlap between patents exceeding the task and tool thresholds changes. Mean patent-text length is essentially unchanged between 1980 and 2021 (760 versus 751 characters), suggesting that the decline is not simply driven by shorter descriptions. Accordingly, these annual category shares are best viewed as descriptive rather than as direct measures of changes in the economy-wide prevalence of labor-automating or labor-augmenting technological change. Figure D.2 reinforces this distinction by showing that the absolute numbers of both automating

and augmenting patents rise substantially over the period despite their declining shares of the overall patent stock.

Third, Figures D.3–D.5 characterize the technological composition of automating and augmenting patents and its evolution over time. Among AI patents, the two classes display broadly similar distributions across the eight AIPD technology functions, with computer vision and planning and decision systems accounting for the largest shares in both channels. Their distributions across Cooperative Patent Classification (CPC) sections are likewise relatively similar. AI patents in both classes are heavily concentrated in Section G (Physics), which contains much of modern computing and information processing. By contrast, non-AI patents are distributed much more broadly across CPC sections over the sample period.

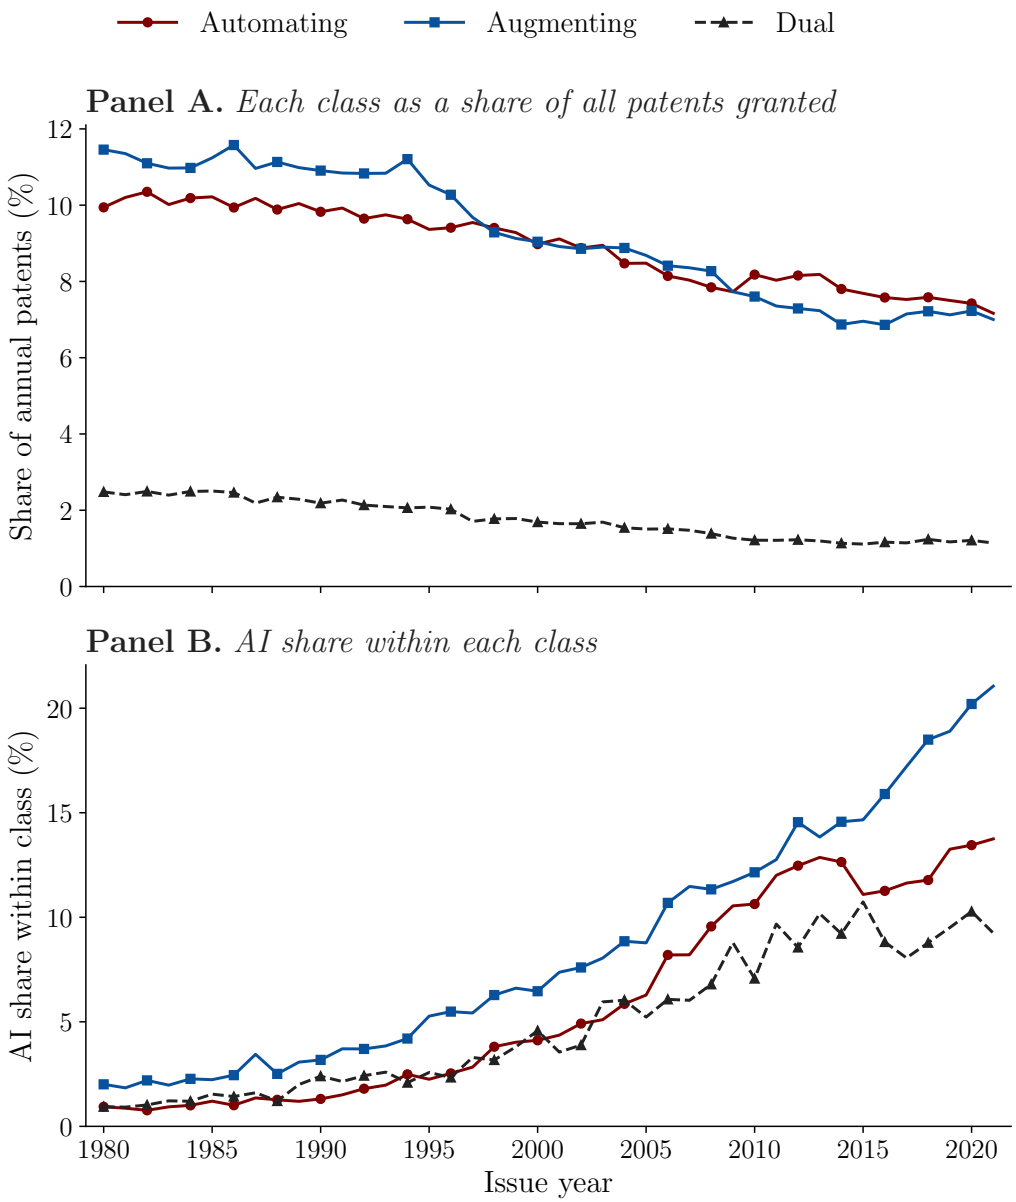


Figure D.1: The Composition of the U.S. Patent Stock: Automation, Augmentation, and AI

Notes: The upper panel reports the annual share of U.S. utility patents classified as automating, augmenting, or dual based on the patent-level classification procedure described in Appendix D.1. Patents classified as neither automating nor augmenting are omitted from the figure. Automating patents exceed only the 90th-percentile threshold of the per-patent maximum task-similarity distribution, augmenting patents exceed only the corresponding threshold of the per-patent maximum tool-similarity distribution, and dual patents exceed both thresholds. The lower panel reports the share of patents classified as AI-related within each patent-level category using the AI classification of Pairolero et al. (2025). Annual shares are calculated using USPTO utility patents granted between 1980 and 2021.

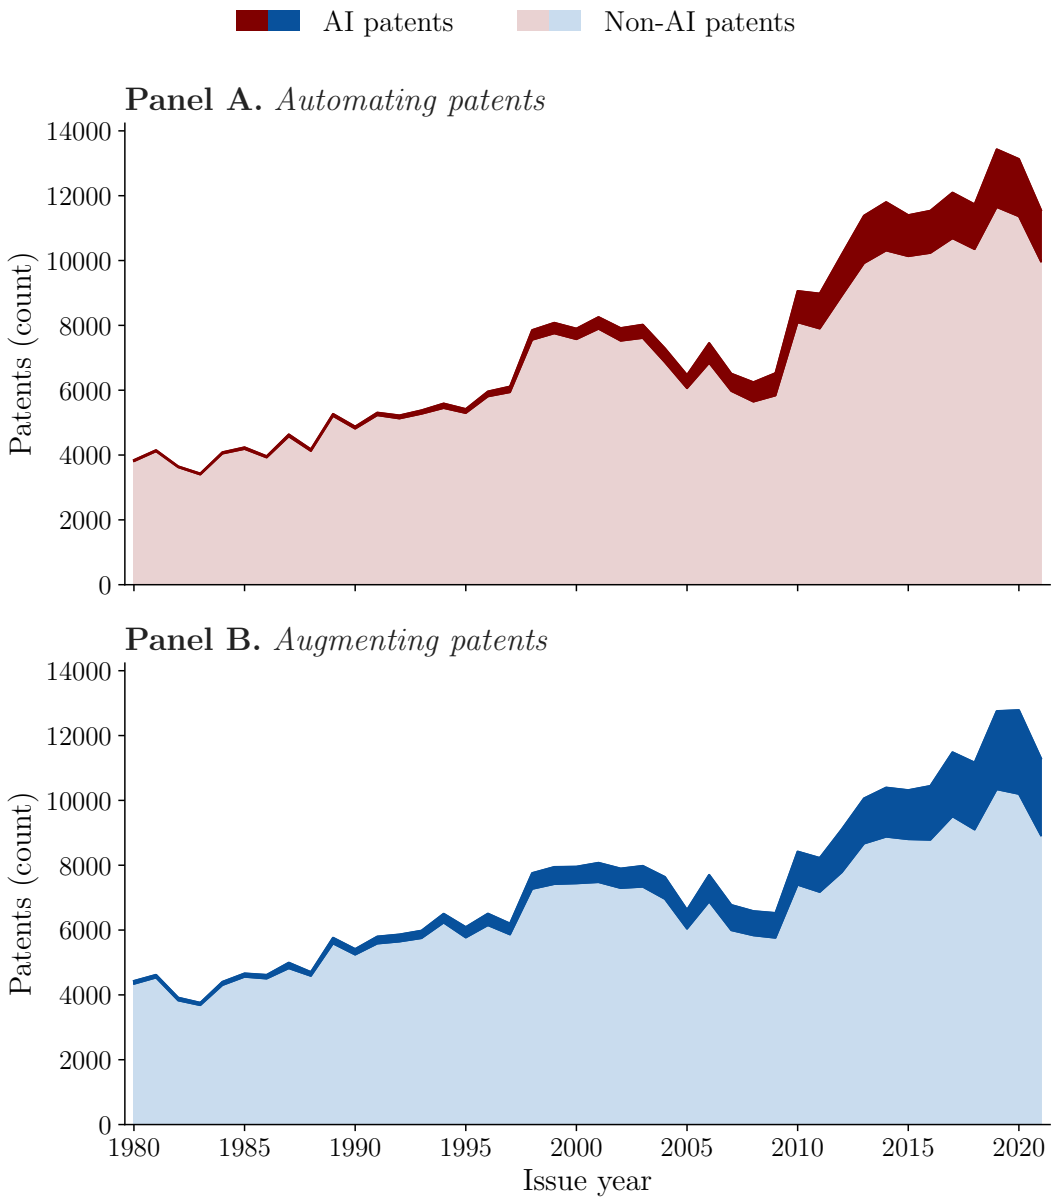


Figure D.2: Annual Counts of Automating and Augmenting Patents by AI Status

Notes: The figure plots the annual number of patents classified as automating (Panel A) and augmenting (Panel B) based on the patent-level classification described in Appendix D.1. Within each class, patents are divided into AI and non-AI patents using the AI classification of Pairolero et al. (2025). The shaded areas are stacked, so their combined height gives the total number of patents in the corresponding class in each year, with the darker area representing AI patents and the lighter area non-AI patents.

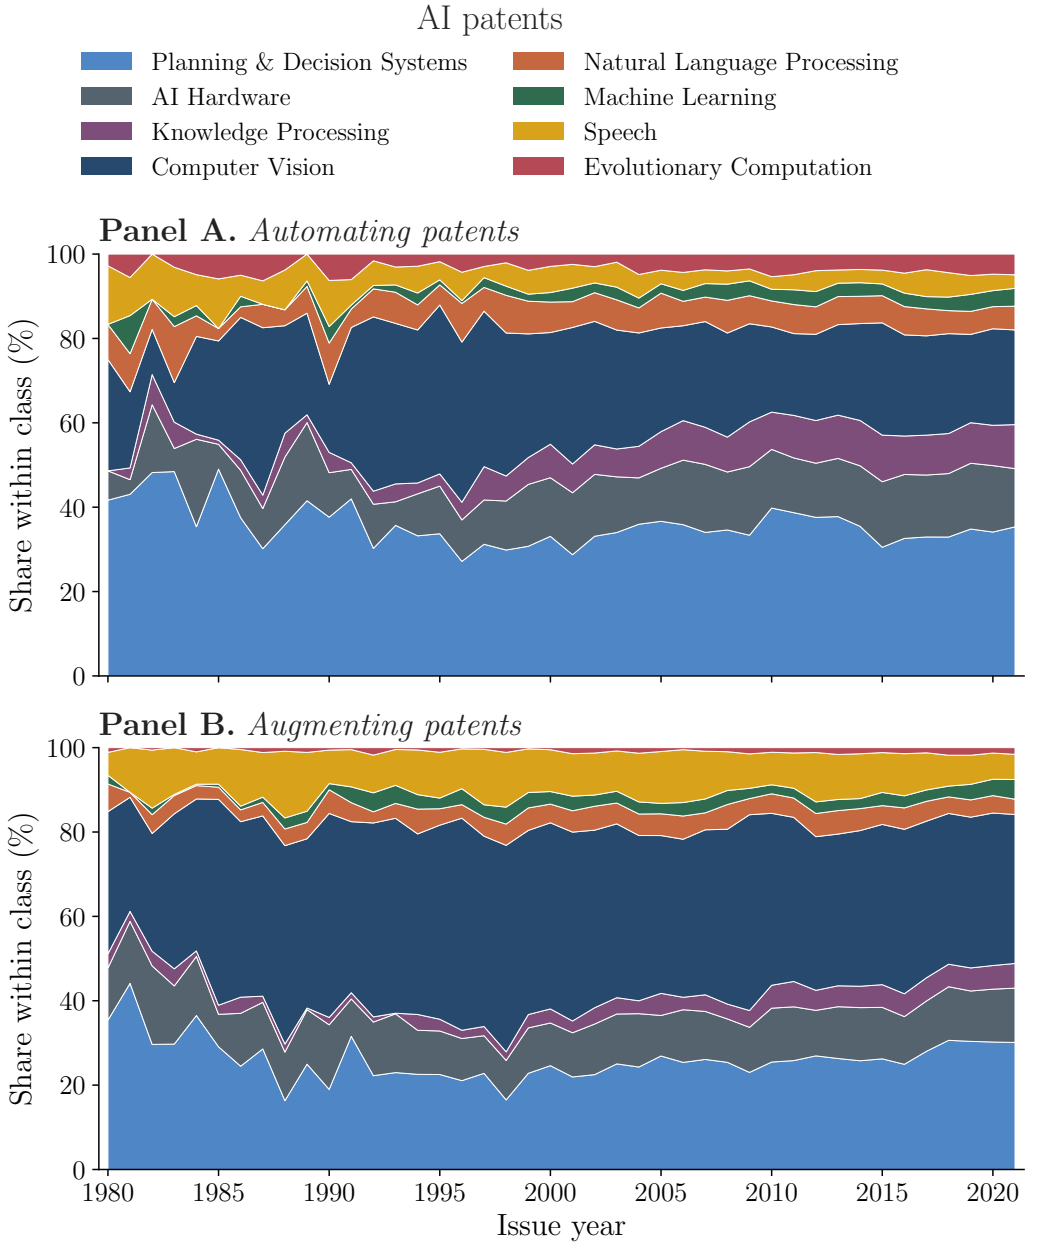


Figure D.3: Evolution of AI Function Composition within Automating and Augmenting Patents

Notes: The figure plots the annual distribution of AI technology functions among AI patents classified as automating (Panel A) and augmenting (Panel B). Within each panel and year, function shares are calculated relative to the corresponding set of AI patents and sum to 100 percent. Patents classified into multiple AI technology functions are fractionally allocated across those functions such that each patent contributes a total weight of one. The patent-level automation and augmentation classification follows the procedure described in Appendix D.1. AI patents and their technology functions are identified using the AI classification of Pairolero et al. (2025).

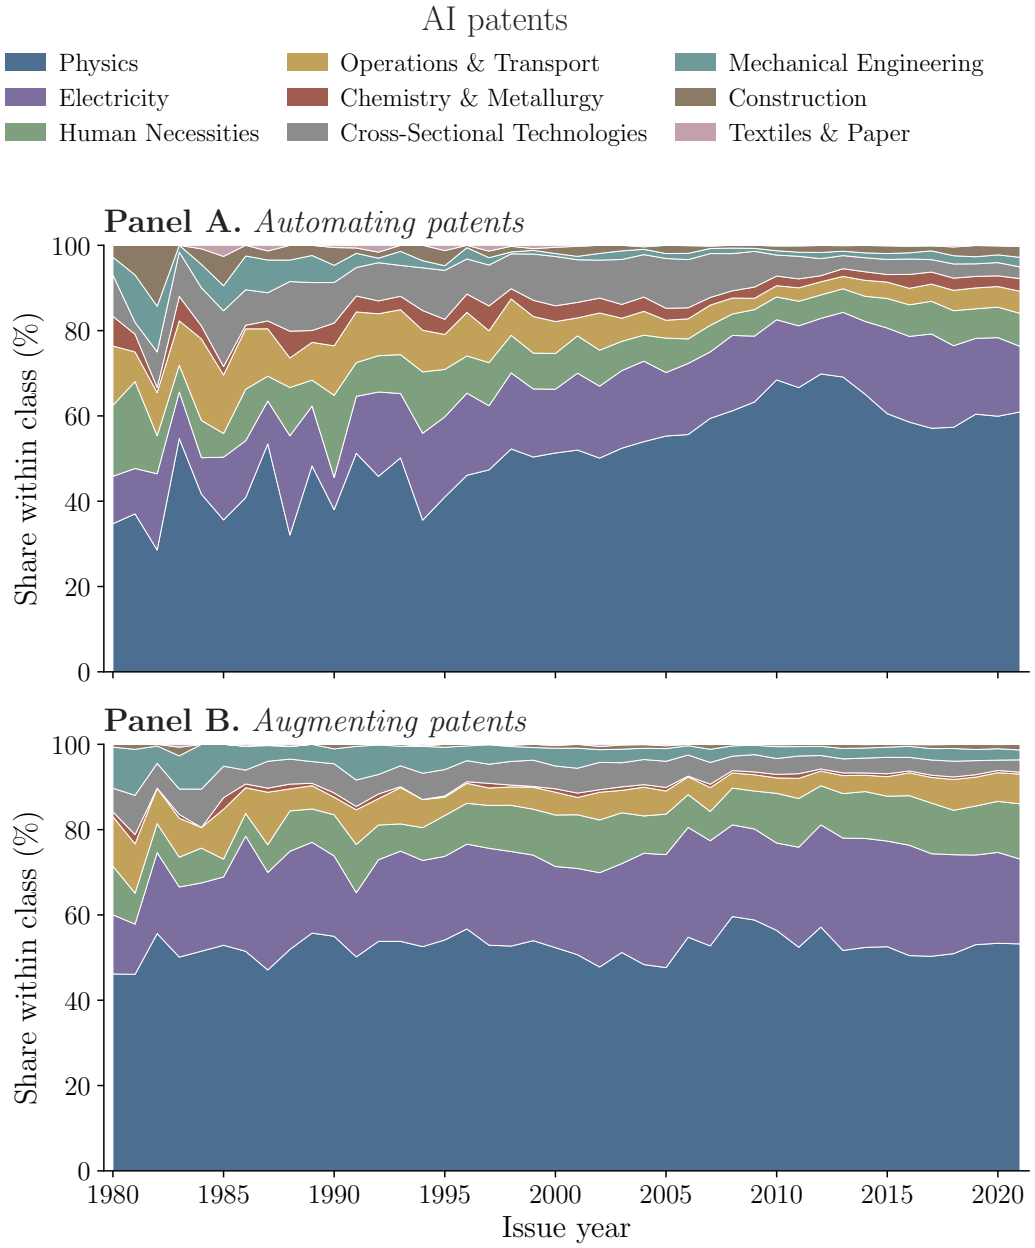


Figure D.4: CPC Technology Composition within Automating and Augmenting AI Patents

Notes: The figure plots the annual distribution of Cooperative Patent Classification (CPC) technology sections among AI patents classified as automating (Panel A) and augmenting (Panel B). Within each panel and year, section shares are calculated relative to the corresponding set of AI patents and sum to 100 percent. Patents assigned to multiple CPC technology sections are fractionally allocated across those sections such that each patent contributes a total weight of one. The patent-level automation and augmentation classification follows the procedure described in Appendix D.1. AI patents are identified using the AI classification of Pairolero et al. (2025).

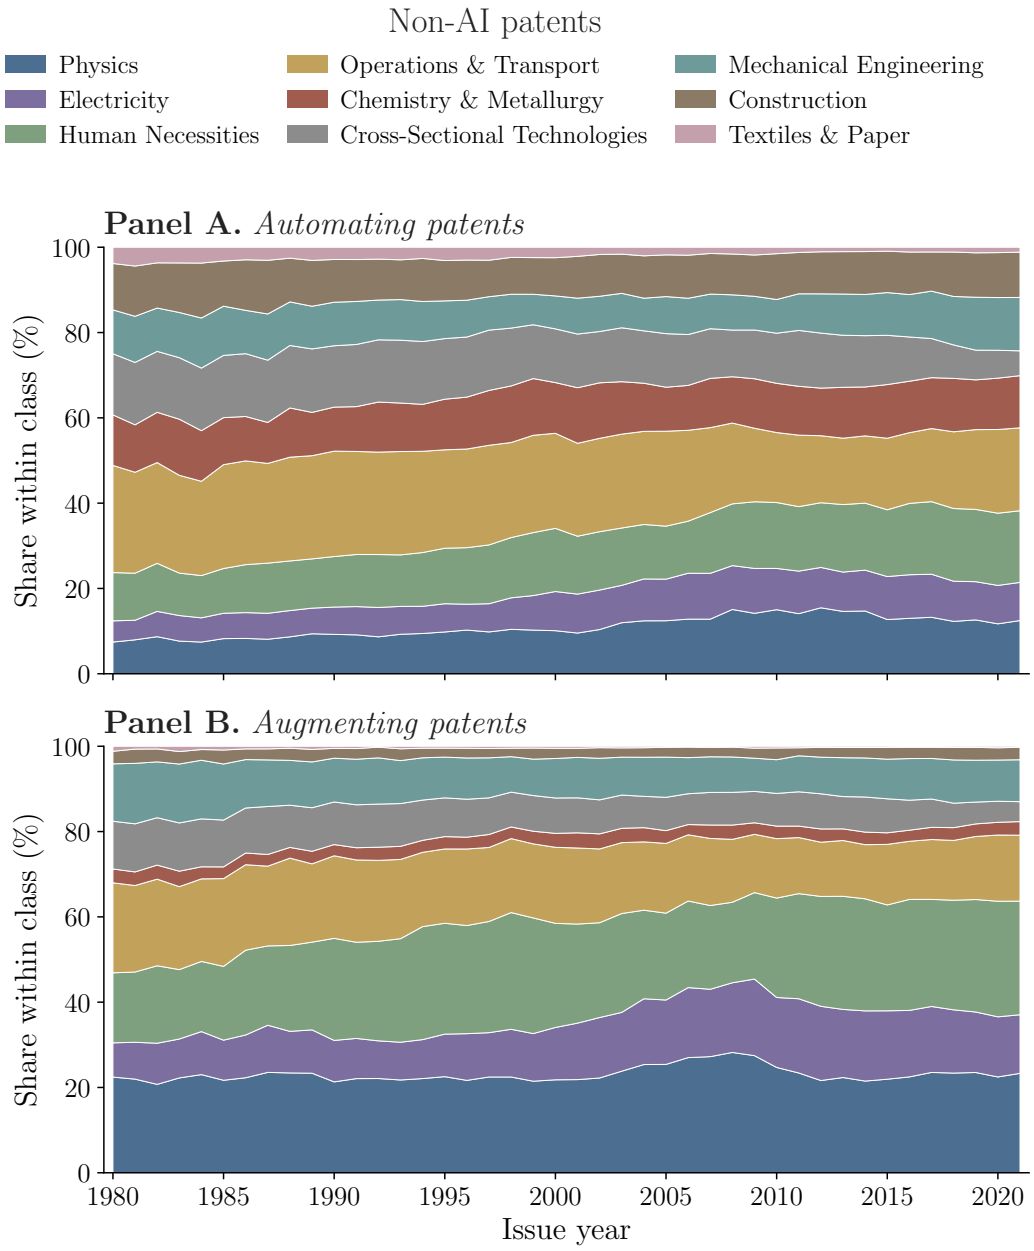


Figure D.5: CPC Technology Composition within Automating and Augmenting Non-AI Patents

Notes: The figure plots the annual distribution of Cooperative Patent Classification (CPC) technology sections among non-AI patents classified as automating (Panel A) and augmenting (Panel B). Within each panel and year, section shares are calculated relative to the corresponding set of non-AI patents and sum to 100 percent. Patents assigned to multiple CPC technology sections are fractionally allocated across those sections such that each patent contributes a total weight of one. The patent-level automation and augmentation classification follows the procedure described in Appendix D.1. Non-AI patents are defined as patents not identified as AI-related using the classification of Pairolero et al. (2025).

E Measurement and Exposure Robustness

Alternative similarity thresholds. We first examine whether the estimated effects depend on the similarity threshold used to define patent–task and patent–tool matches, considering alternative cutoffs from the 75th to the 99th percentile. Table E.1 shows that the negative effects of AI automation on employment and the wage bill remain statistically significant across all thresholds, while its hourly-wage effect is more sensitive to the cutoff but remains comparatively small throughout. The estimated wage-bill effect of AI augmentation is positive at every cutoff and statistically significant from the 85th percentile onward. For AI augmentation, the underlying employment and wage responses display a systematic pattern: as the matching criterion becomes more stringent, the estimated employment effect shifts from negative to increasingly positive, whereas the positive wage effect becomes smaller. However, even as its magnitude declines, the positive wage response remains statistically significant at every cutoff.

Alternative augmentation measures. Table E.2 assesses the robustness of the augmentation results to two choices in the construction of the exposure measure: the method used to weight the relevance of occupational tools and the language model used to match patents to tools. First, weighting tools according to their similarity to occupation-title profiles constructed from the O*NET Alternate Titles database, rather than to occupations' task profiles, produces estimates that are very close to the baseline results. Second, we reconstruct AI augmentation exposure and its corresponding historical instrument using the general-purpose MPNet model in place of PatentSBERTa. While the patent-domain specialization of PatentSBERTa is a deliberate feature of our baseline design, the MPNet specification tests to what extent the results depend on this model choice. We leave the AI automation exposure measure and its corresponding instrument unchanged, retaining JobBERT V2 for patent–task matching. Under the MPNet specification, the wage-bill and hourly-wage effects of AI augmentation remain positive and statistically significant, while the positive employment effect becomes smaller and statistically insignificant.

Automation and augmentation across patent universes. To better understand how the automation–augmentation distinction differs between AI and non-AI innovation, we retain the baseline workforce controls and commuting-zone-by-broad-occupation fixed effects and construct the exposure measures separately from AI patents, the full patent universe, and non-AI patents. In each case, automation and augmentation exposure are jointly instrumented using the corresponding historical breakthrough exposure measures. Table E.3 shows that the estimates based on the full patent universe generally lie between those obtained for AI and non-AI patents. Comparing the two patent groups reveals two additional patterns consistent with the results for the relative AI exposure measures in Section 6.3. First, the negative employment and wage-bill effects of automation are substantially larger in magnitude for AI than for non-AI exposure, with the employment point estimate for AI automation nearly twice as large. Second, augmentation operates through different adjustment margins: non-AI augmentation yields a positive and statistically significant employment effect, whereas the employment effect of AI augmentation is

small and statistically insignificant. By contrast, the positive wage effect of AI augmentation is markedly larger than that of non-AI augmentation.

Exposure from neighboring commuting zones. The baseline exposure measures are constructed using patents originating in the focal commuting zone. To assess sensitivity to this local definition and allow exposure to innovative activity to extend across commuting-zone boundaries, we expand contemporary exposure to include patents originating in immediately adjacent commuting zones, while leaving the historical instruments unchanged. Table E.4 shows that the main patterns remain largely unchanged: AI automation continues to reduce employment and the wage bill, while AI augmentation raises hourly wages and the wage bill and yields positive employment estimates. The Kleibergen–Paap F-statistic also increases substantially under this broader spatial definition.

Positive-exposure sample. Because the exposure measures are constructed from local patenting, a substantial number of occupation–commuting-zone cells have zero AI automation or augmentation exposure. To assess whether the baseline results are driven by comparisons involving these cells, we restrict the sample to cells with strictly positive local AI exposure in both channels. We therefore rely exclusively on intensive-margin variation among cells exposed to both automation and augmentation. Table E.5 shows that the estimates remain similar in sign and magnitude to the baseline results.

Relative AI exposure measures. We conduct analogous robustness exercises for the relative AI exposure measures, which capture the AI share of total automation and augmentation exposure within each channel. Table E.6 indicates that the main composition effects are generally stable across alternative similarity thresholds, augmentation measures, spatial exposure definitions, and the positive-exposure sample. A higher AI share within automation produces negative employment and wage-bill estimates across the alternative constructions, except under the MPNet specification, where the wage-bill estimate turns positive and the employment effect is close to zero. The hourly-wage effect of relative AI automation remains comparatively small. For augmentation, the central pattern of lower employment growth and higher hourly-wage growth is preserved across all alternative constructions, although the MPNet estimates are more sensitive in magnitude and imply a negative wage-bill effect.

Table E.1: Robustness to Alternative Patent–Task and Patent–Tool Similarity Thresholds

	Percentile cutoff				
	75th	85th	**90th**	95th	99th
Panel A: Wage bill change					
AI Automation	-0.249***	-0.257***	-0.248***	-0.205***	-0.164***
	(0.035)	(0.023)	(0.019)	(0.016)	(0.015)
AI Augmentation	0.044	0.116***	0.145***	0.146***	0.134***
	(0.036)	(0.029)	(0.026)	(0.023)	(0.022)
Panel B: Employment change					
AI Automation	-0.191***	-0.228***	-0.231***	-0.203***	-0.181***
	(0.034)	(0.021)	(0.016)	(0.013)	(0.014)
AI Augmentation	-0.109***	0.013	0.067***	0.101***	0.140***
	(0.029)	(0.023)	(0.020)	(0.018)	(0.019)
Panel C: Hourly wage change					
AI Automation	-0.029**	-0.017**	-0.011	-0.002	0.011**
	(0.012)	(0.009)	(0.007)	(0.006)	(0.005)
AI Augmentation	0.139***	0.100***	0.079***	0.056***	0.016**
	(0.014)	(0.011)	(0.010)	(0.008)	(0.007)
KP F-statistic	691.03	698.60	711.01	746.70	520.56
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using the common estimation sample of column (5) and the preferred specification in column (4) of Table 3. The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages between 2005–07 and 2022–24. Columns vary the percentile threshold used to define patent–task and patent–tool matches; the 90th percentile is the baseline threshold. At each threshold, AI automation and augmentation exposure, the corresponding historical breakthrough instruments, and the non-AI automation and augmentation exposure controls are constructed using the same cutoff. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within the common estimation sample. All specifications include the corresponding threshold-specific non-AI exposure controls, baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), and commuting-zone-by-broad-occupation fixed effects. Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table E.2: Robustness to Alternative Measures of AI Augmentation

	Baseline	O*NET Alternate Titles weights	MPNet model
Panel A: Wage bill change			
AI Automation	-0.248***	-0.247***	-0.222***
	(0.019)	(0.019)	(0.019)
AI Augmentation	0.145***	0.141***	0.072***
	(0.026)	(0.025)	(0.025)
Panel B: Employment change			
AI Automation	-0.231***	-0.229***	-0.211***
	(0.016)	(0.016)	(0.017)
AI Augmentation	0.067***	0.063***	0.017
	(0.020)	(0.020)	(0.020)
Panel C: Hourly wage change			
AI Automation	-0.011	-0.012*	-0.005
	(0.007)	(0.007)	(0.007)
AI Augmentation	0.079***	0.080***	0.056***
	(0.010)	(0.010)	(0.009)
KP F-statistic	711.01	711.51	699.37
Observations	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using the common estimation sample of column (5) and the preferred specification in column (4) of Table 3. Column (1) reports the baseline estimates. Column (2) replaces the baseline AI augmentation exposure measure and corresponding historical instrument with measures using occupation-title-profile tool weights constructed from O*NET Alternate Titles. Column (3) replaces the baseline AI augmentation exposure measure and corresponding historical instrument with MPNet-based measures. The baseline AI automation exposure measure and corresponding historical instrument are retained throughout. AI exposure measures are inverse-hyperbolic-sine transformed and standardized within the common estimation sample. All specifications include baseline non-AI automation and augmentation controls, baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), and commuting-zone-by-broad-occupation fixed effects. Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table E.3: Automation and Augmentation across Patent Universes

	(1) AI patents	**(2)** All patents	**(3)** Non-AI patents
Panel A: Wage bill change			
Automation	-0.105***	-0.067***	-0.061***
	(0.013)	(0.013)	(0.013)
Augmentation	0.030*	0.097***	0.124***
	(0.016)	(0.024)	(0.021)
Panel B: Employment change			
Automation	-0.109***	-0.069***	-0.060***
	(0.011)	(0.011)	(0.011)
Augmentation	-0.022	0.035	0.083***
	(0.015)	(0.026)	(0.022)
Panel C: Hourly wage change			
Automation	-0.002	-0.014***	-0.018***
	(0.003)	(0.004)	(0.004)
Augmentation	0.060***	0.050***	0.025***
	(0.006)	(0.008)	(0.007)
KP F-statistic	722.96	264.35	420.40
Observations	93,026	93,026	93,026

Notes: The table compares 2SLS estimates of automation and augmentation exposure constructed from alternative patent universes using the common estimation sample of column (5) of Table 3. Column (1) uses AI patents, column (2) the full patent universe, and column (3) non-AI patents. In each column, automation and augmentation are jointly instrumented using the corresponding historical breakthrough exposure measures constructed from the same patent universe. Contemporary exposure measures are inverse-hyperbolic-sine transformed and standardized within the common estimation sample. All specifications include baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing) and commuting-zone-by-broad-occupation fixed effects. Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table E.4: Robustness to AI Exposure from Own and Neighboring Commuting Zones

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.136***	-0.123***	-0.265***	-0.272***	-0.130***
	(0.031)	(0.027)	(0.020)	(0.019)	(0.022)
AI Augmentation	0.172***	0.151***	0.164***	0.146***	0.099***
	(0.042)	(0.040)	(0.027)	(0.027)	(0.025)
Panel B: Employment change					
AI Automation	-0.114***	-0.102***	-0.240***	-0.246***	-0.153***
	(0.026)	(0.022)	(0.017)	(0.016)	(0.017)
AI Augmentation	0.115***	0.075**	0.082***	0.063***	0.032
	(0.035)	(0.032)	(0.021)	(0.021)	(0.020)
Panel C: Hourly wage change					
AI Automation	-0.025***	-0.025***	-0.021***	-0.021***	0.003
	(0.007)	(0.007)	(0.008)	(0.008)	(0.008)
AI Augmentation	0.055***	0.074***	0.089***	0.087***	0.079***
	(0.012)	(0.011)	(0.010)	(0.010)	(0.010)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	575.62	1484.03	1271.45	1667.12	1660.32
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells after expanding contemporary AI automation and augmentation exposure to include patents originating in the focal commuting zone and directly adjacent commuting zones. All specifications use the common estimation sample of column (5) of Table 3. The historical instruments are unchanged from the baseline specification. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within the common estimation sample. All specifications control for non-AI automation and augmentation exposure constructed using the same own-plus-adjacent-CZ definition and baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table E.5: Intensive-Margin Robustness: Positive Local AI Exposure in Both Channels

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.165***	-0.158***	-0.268***	-0.262***	-0.112***
	(0.034)	(0.029)	(0.021)	(0.020)	(0.022)
AI Augmentation	0.184***	0.229***	0.183***	0.158***	0.107***
	(0.048)	(0.047)	(0.027)	(0.026)	(0.024)
Panel B: Employment change					
AI Automation	-0.157***	-0.139***	-0.248***	-0.243***	-0.142***
	(0.029)	(0.025)	(0.019)	(0.017)	(0.018)
AI Augmentation	0.141***	0.149***	0.102***	0.078***	0.042**
	(0.040)	(0.039)	(0.021)	(0.020)	(0.019)
Panel C: Hourly wage change					
AI Automation	-0.014*	-0.020***	-0.016*	-0.012*	0.011
	(0.008)	(0.007)	(0.008)	(0.007)	(0.008)
AI Augmentation	0.044***	0.071***	0.084***	0.080***	0.073***
	(0.013)	(0.012)	(0.010)	(0.010)	(0.010)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	602.55	726.61	620.22	798.99	791.13
Observations	69,238	69,238	69,238	69,238	69,238

Notes: The sample is restricted to occupation–commuting-zone cells with strictly positive local AI automation and augmentation exposure during 2005–2021. The restriction is imposed on the common estimation sample of column (5) of Table 3, and all columns and panels use the resulting common restricted sample. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table E.6: Robustness of Relative AI Automation and Augmentation Effects

		Similarity thresholds		*Augmentation measures*		*Exposure and sample*	
	Baseline	p85	p95	Alt. title weights	MPNet model	Adjacent CZs	Positive exposure
Panel A: Wage bill change							
Relative AI Automation	-0.030***	-0.030***	-0.016**	-0.027***	0.016**	-0.029***	-0.030***
	(0.008)	(0.009)	(0.007)	(0.008)	(0.008)	(0.007)	(0.008)
Relative AI Augmentation	0.009	0.003	0.004	0.003	-0.078***	0.003	0.009
	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)	(0.006)	(0.007)
Panel B: Employment change							
Relative AI Automation	-0.027***	-0.026***	-0.010	-0.022***	-0.001	-0.025***	-0.029***
	(0.008)	(0.010)	(0.007)	(0.008)	(0.009)	(0.008)	(0.009)
Relative AI Augmentation	-0.026***	-0.035***	-0.031***	-0.033***	-0.078***	-0.026***	-0.024***
	(0.007)	(0.008)	(0.007)	(0.007)	(0.008)	(0.006)	(0.007)
Panel C: Hourly wage change							
Relative AI Automation	0.005*	0.006**	0.004	0.005*	0.019***	0.005*	0.007**
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Relative AI Augmentation	0.029***	0.031***	0.028***	0.030***	0.006*	0.024***	0.028***
	(0.003)	(0.004)	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
KP F-statistic	464.64	341.33	609.59	418.17	552.17	537.70	510.57
Observations	86,645	86,645	86,645	86,645	86,645	86,645	68,489

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using the preferred specification in column (4) of Table 5. The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages between 2005–07 and 2022–24. Relative AI automation and augmentation exposure are defined as the AI shares of total exposure within each channel and instrumented using the corresponding AI shares of total historical-breakthrough exposure. Column (1) uses the baseline 90th-percentile construction. Columns (2)–(3) use the 85th- and 95th-percentile similarity thresholds, with the corresponding threshold-specific total-exposure controls. Columns (4)–(5) replace the baseline relative AI augmentation measure and corresponding historical instrument with measures based on occupation-title-profile tool weights constructed from O*NET Alternate Titles and the MPNet matching model, respectively, while retaining the baseline relative AI automation measure and its corresponding historical instrument, as well as the baseline total automation and augmentation controls. Column (6) expands contemporary exposure to include the focal commuting zone and directly adjacent commuting zones, uses total-exposure controls constructed with the same own-plus-adjacent-CZ definition, and retains the baseline historical instruments. Columns (1)–(6) use a common robustness sample nested within the common estimation sample of column (5) of Table 5 and requiring all alternative relative measures and corresponding instruments to be defined. Column (7) further restricts this sample to cells with strictly positive baseline local AI automation and augmentation exposure. Relative AI exposure measures are standardized within each estimation sample. All specifications include baseline workforce controls (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing) and commuting-zone-by-broad-occupation fixed effects. Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

F Supporting Evidence for the Baseline IV Design

This section documents the historical and spatial structure underlying the instruments, assesses whether historical breakthroughs generate sufficiently distinct variation in automation and augmentation exposure, and examines reduced-form relationships and inference robustness. Figure F.1 shows the evolution of AI and AI breakthrough patenting and marks the breakthrough period used to construct the instruments and the contemporaneous patent period used to measure local AI exposure. During the instrument period, the AI share of breakthrough patents rose from around 5 percent in 1980 to roughly 25 percent by 2000. Over the same period, the AI share of all patents increased more gradually and remained below 10 percent.

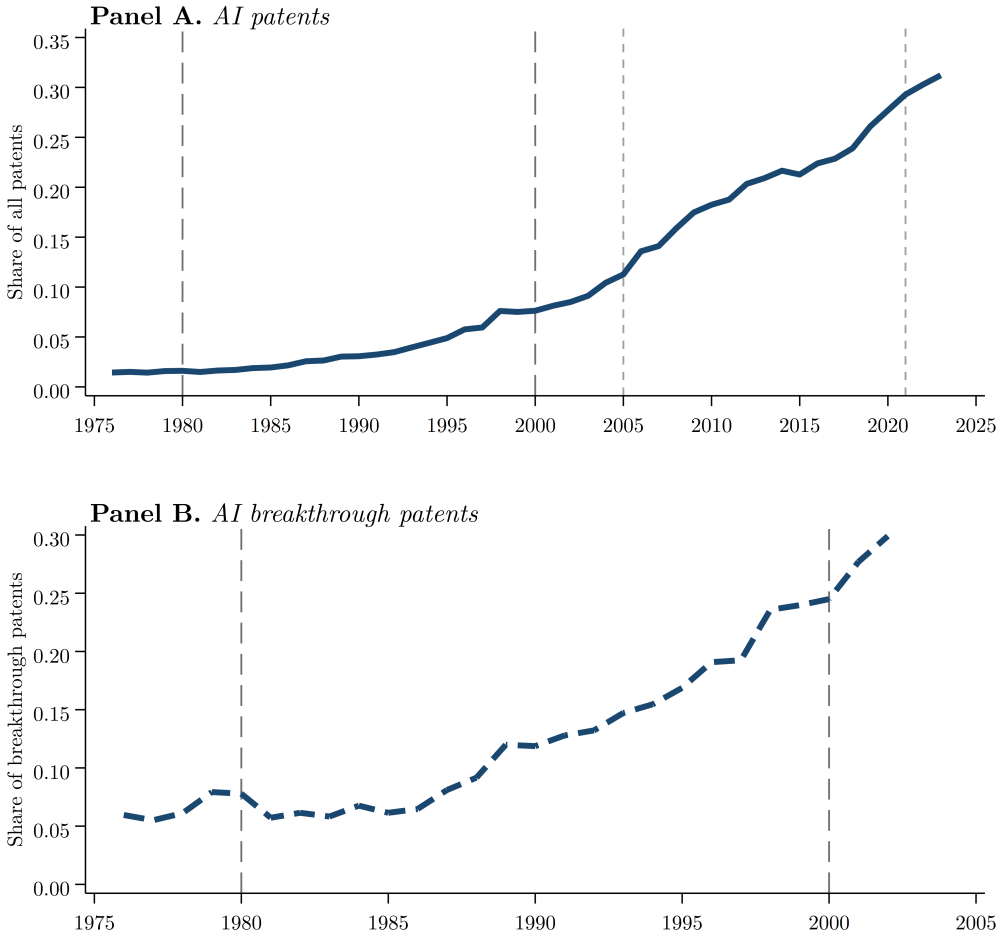


Figure F.1: Evolution of AI Patenting and AI Breakthrough Patenting in the United States

Notes: Panel A plots the annual share of AI patents among geographically assigned domestic U.S. utility patents, and Panel B the AI share among breakthrough patents. AI patents are identified using the USPTO Artificial Intelligence Patent Dataset (AIPD), classifying a patent as AI-related when its predicted probability for at least one AI function is at least 86 percent (Pairolero et al. 2025). Breakthrough patents are defined as patents in the top decile of the grant-year distribution of the residual forward-similarity measure of Kelly et al. (2021). Vertical lines mark the 1980–2000 historical breakthrough window and the 2005–2021 contemporaneous exposure window.

Two-Ring Exclusion Buffer. To characterize the spatial correlation in contemporaneous AI exposure, let $E^k_{oc}(d)$ denote the raw exposure of occupation o in focal commuting zone c originating from patents in commuting zones located exactly d contiguity rings away. The index $k \in \{\text{Aut}, \text{Aug}\}$ denotes the automation or augmentation channel, and $E^k_{oc}(0)$ denotes exposure originating in the focal commuting zone. For each distance d, Appendix Figure F.2 reports

$$\rho^k(d) = \text{Corr}_{(o,c)\in\mathcal{S}_d}\left(E^k_{oc}(0), E^k_{oc}(d)\right), \tag{27}$$

where $\mathcal{S}_d$ contains occupation–commuting-zone cells for which the focal commuting zone has at least one donor commuting zone at that distance. In both channels, these unweighted correlations decline sharply across the nearest rings and remain close to zero at greater distances. This concentration of spatial co-movement within the nearest rings is consistent with localized spatial dependence in innovative activity and motivates excluding the focal commuting zone and its first two surrounding rings when constructing the instruments.

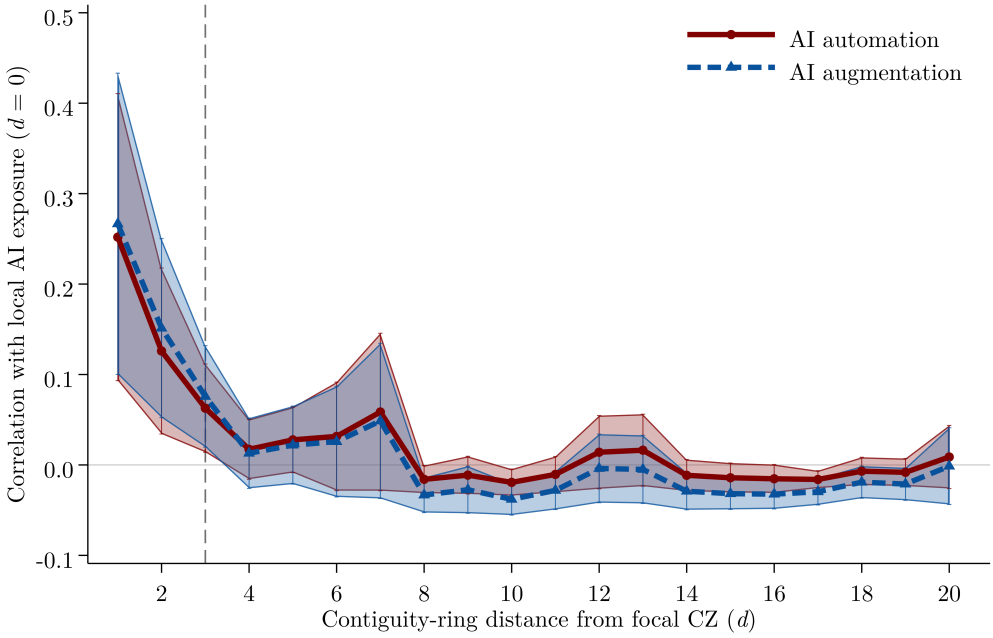


Figure F.2: Correlation Between Local and Ring-Specific AI Exposure

Notes: The figure reports unweighted correlations across occupation–commuting-zone cells between raw local AI exposure and exposure originating from AI patents granted between 2005 and 2021 in commuting zones located exactly d contiguity rings from the focal commuting zone. Rings 1–20 are shown; at each distance, focal commuting zones whose contiguity networks do not extend to that ring are omitted. Shaded bands report 95 percent confidence intervals based on 1,000 bootstrap replications clustered by focal commuting zone. The dashed vertical line separates the two excluded rings from ring 3, the first ring included in the instruments.

Geographic Distance and Historical Breakthrough Exposure. The construction of our spatial instruments is motivated by the idea that historical breakthroughs closer to a focal commuting zone may contain more predictive information about subsequent local AI exposure than

more distant breakthroughs. We therefore examine how the predictive content of historical AI breakthroughs varies with distance. This exercise complements Appendix Figure F.2, which motivates the two-ring exclusion buffer using contemporaneous exposure, and provides descriptive evidence on the spatial structure of historical breakthrough exposure beyond ring 2. We group donor commuting zones into five mutually exclusive distance bands based on contiguity-ring distance from the focal commuting zone: the focal CZ, rings 1–2, rings 3–5, rings 6–10, and rings 11 and beyond.

Let $B \in \{0, 1\text{–}2, 3\text{–}5, 6\text{–}10, 11+\}$ denote mutually exclusive distance bands, where $B = 0$ refers to the focal commuting zone. For channel $k \in \{\text{Aut}, \text{Aug}\}$, we define average historical breakthrough exposure from band B as

$$\overline{\text{BT}}^k_{oc}(B) = \frac{\sum_{j:d_{cj}\in B} \text{BT}^k_{oj}}{N_c(B)}, \tag{28}$$

where BT^k_{oj} denotes occupation-specific exposure in channel k to historical AI breakthrough patents originating in commuting zone j, and $N_c(B)$ is the number of donor commuting zones in band B. Dividing by the number of donors avoids mechanically larger exposure in more distant bands simply because they contain more commuting zones.

For each band, we calculate

$$\rho^k(B) = \text{Corr}_w\left(\text{asinh}\left(E^k_{oc}(0)\right), \text{asinh}\left(\overline{\text{BT}}^k_{oc}(B)\right)\right). \tag{29}$$

where $E^k_{oc}(0)$ is the contemporary local exposure measure used in the main analysis. Correlations are weighted by 2005–07 FTE employment shares and estimated on a common sample of occupation–commuting-zone cells nested within the common estimation sample of column (5) of Table 3 and for which both contemporary local AI exposure measures and historical breakthrough exposure are observed in all five bands. For inference, both variables are standardized using the same weights, and the corresponding weighted bivariate regressions are estimated with standard errors clustered by commuting zone.

Appendix Table F.1 shows a clear geographic gradient in both channels. For automation, the weighted correlation declines from 0.925 in the focal commuting zone to 0.479 in rings 1–2, 0.338 in rings 3–5, and 0.124 in rings 6–10. For augmentation, the corresponding correlations decline from 0.942 to 0.482, 0.351, and 0.100. Adjacent-band Wald tests indicate statistically significant declines at the 10 percent level through rings 6–10 in both channels. Beyond ring 10, the correlations do not differ significantly from those in rings 6–10 ($p = 0.361$ for automation and $p = 0.828$ for augmentation). This pattern indicates that the predictive content of historical breakthroughs is substantially stronger among geographically closer donor commuting zones and broadly levels off at greater distances.

It is important to note that this geographic pattern is not required for identification and should not be interpreted as a test of the exclusion restriction or of the exact inverse-distance functional form. Rather, it provides descriptive support for assigning greater weight to nearer

Table F.1: Geographic Distance and the Predictive Content of Historical AI Innovation

	Own CZ	Rings 1–2	Rings 3–5	Rings 6–10	Rings 11+
Panel A: Local AI automation exposure					
Weighted correlation	0.925	0.479	0.338	0.124	0.185
	(0.025)	(0.085)	(0.059)	(0.045)	(0.041)
Equality with preceding band (p-value)	–	< 0.001	0.052	0.004	0.361
Panel B: Local AI augmentation exposure					
Weighted correlation	0.942	0.482	0.351	0.100	0.117
	(0.022)	(0.086)	(0.073)	(0.052)	(0.044)
Equality with preceding band (p-value)	–	< 0.001	0.075	0.004	0.828
Observations	91,316	91,316	91,316	91,316	91,316

Notes: The table reports weighted correlations between contemporary local AI exposure and historical AI breakthrough exposure originating from five mutually exclusive contiguity-distance bands: the focal commuting zone, rings 1–2, rings 3–5, rings 6–10, and rings 11 and beyond. Historical exposure is measured per donor commuting zone, and both contemporary and historical exposure measures are inverse-hyperbolic-sine transformed. All correlations are estimated on a common sample of occupation–commuting-zone cells nested within the common estimation sample of column (5) of Table 3 and for which both contemporary local AI exposure measures and historical breakthrough exposure are observed in all five bands. Estimates are weighted by 2005–07 FTE employment shares. For inference, the two variables entering each correlation are standardized using the same weights, and the corresponding weighted bivariate regressions are estimated with standard errors clustered by commuting zone, reported in parentheses. The Wald-test rows report cluster-robust p-values for equality of each correlation with the correlation in the immediately preceding distance band.

eligible donor commuting zones. Complementing this exercise, Appendix Table H.4 shows that the estimates are robust to using steeper inverse-square distance weights.

Correlation Between AI Automation and Augmentation Exposure. Appendix Figure F.3 visualizes the relationship between the automation and augmentation measures of AI exposure. Panel A shows that the contemporary local exposure measures, $E_{oc}^{\text{Aut}}(0)$ and $E_{oc}^{\text{Aug}}(0)$, are strongly positively correlated (corr = 0.95), reflecting their shared geographic origins and the fact that the same AI patents may match both occupational tasks and tools. By contrast, Panel B shows that the corresponding historical-breakthrough instruments are clearly less correlated (corr = 0.48). The instruments therefore contain substantially more distinct variation across the two channels. Together with the conditional first-stage statistics reported in Table 2, this provides supporting evidence that the instruments retain sufficient predictive power for each exposure measure conditional on the other, allowing us to distinguish the two channels empirically.

Reduced-Form Evidence. Appendix Table F.2 reports reduced-form estimates relating the historical AI automation and augmentation instruments directly to the three main labor-market outcomes, with both instruments entered jointly. Across all specifications, the automation instrument is negatively associated with wage-bill and employment growth, while its association with hourly-wage growth is comparatively small and less stable. The augmentation instrument is positively associated with wage-bill and hourly-wage growth throughout. Its association with employment growth is positive and statistically significant in columns (1)–(2) but becomes close

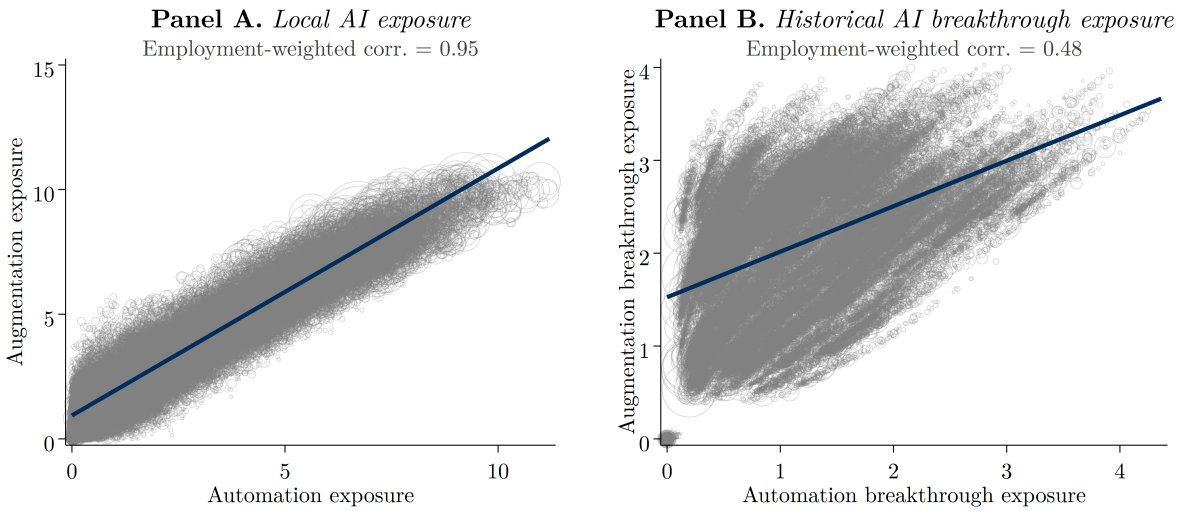


Figure F.3: Correlation Between AI Automation and Augmentation Exposure

Notes: Panel A plots local AI automation and augmentation exposure constructed from AI patents granted between 2005 and 2021. Panel B plots the corresponding instrumental-variable measures constructed from AI breakthrough patents granted between 1980 and 2000 originating outside the two-ring exclusion buffer. Both panels display inverse-hyperbolic-sine-transformed measures. Observations are occupation–commuting-zone cells from the common estimation sample of column (5) of Table 3. Marker area is proportional to 2005–07 FTE employment shares, which are also used to weight the reported correlations and linear fits.

to zero once broad occupation fixed effects are introduced. Although consistent with our main results, these reduced-form coefficients should be interpreted in conjunction with the joint first-stage relationships reported in Table 2.

Inference Robustness. Because the identifying variation is geographically structured around focal commuting zones and occupation-specific outcomes within the same local labor market may be affected by common shocks, we cluster standard errors by commuting zone in our baseline specification. This choice is consistent with the design-based motivation discussed by Abadie et al. (2023), which emphasizes that the clustering level should reflect the structure of the underlying identifying variation. Appendix Table F.3 examines sensitivity to alternative forms of residual dependence. Conley spatial-HAC standard errors (Conley 1999), which allow for residual correlation across nearby commuting zones over a range of distance cutoffs, are similar in magnitude to the baseline commuting-zone-clustered standard errors and leave all statistical conclusions unchanged. We also consider a more conservative specification with two-way clustering by commuting zone and detailed occupation, which additionally allows arbitrary residual correlation within detailed occupations across commuting zones. The resulting standard errors are substantially larger, yielding less precise inference under two-way clustering.

Table F.2: Reduced-Form Estimates for AI Automation and Augmentation Instruments

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
Automation IV	-0.034***	-0.035***	-0.055***	-0.058***	-0.023***
	(0.007)	(0.006)	(0.005)	(0.005)	(0.005)
Augmentation IV	0.025**	0.041***	0.020***	0.017***	0.015***
	(0.010)	(0.010)	(0.006)	(0.006)	(0.005)
Panel B: Employment change					
Automation IV	-0.033***	-0.032***	-0.051***	-0.054***	-0.031***
	(0.006)	(0.006)	(0.004)	(0.004)	(0.004)
Augmentation IV	0.017**	0.023***	0.002	-0.002	-0.003
	(0.008)	(0.008)	(0.005)	(0.005)	(0.004)
Panel C: Hourly wage change					
Automation IV	-0.002	-0.004**	-0.003	-0.003*	0.003*
	(0.002)	(0.002)	(0.002)	(0.002)	(0.002)
Augmentation IV	0.008***	0.016***	0.020***	0.020***	0.020***
	(0.003)	(0.003)	(0.002)	(0.002)	(0.002)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports reduced-form regressions of wage-bill, employment, and hourly-wage changes on the excluded historical AI automation and augmentation instruments, which enter jointly in every specification. All columns and panels use the common estimation sample of column (5) of Table 3. The instruments are inverse-hyperbolic-sine transformed and standardized within this common sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table F.3: Inference Robustness for AI Automation and Augmentation Estimates

	Wage-bill change		***Employment change***		***Hourly-wage change***	
	Automation **(1)**	Augmentation **(2)**	Automation **(3)**	Augmentation **(4)**	Automation **(5)**	Augmentation **(6)**
Point estimate	-0.248	0.145	-0.231	0.067	-0.011	0.079
Standard errors						
CZ cluster	(0.019)	(0.026)	(0.016)	(0.020)	(0.007)	(0.010)
Conley, 500 km	(0.021)	(0.023)	(0.017)	(0.019)	(0.009)	(0.011)
Conley, 1,000 km	(0.021)	(0.022)	(0.018)	(0.020)	(0.009)	(0.011)
Conley, 2,000 km	(0.023)	(0.022)	(0.019)	(0.021)	(0.009)	(0.011)
Conley, 3,500 km	(0.022)	(0.021)	(0.019)	(0.019)	(0.007)	(0.010)
Two-way cluster	(0.193)	(0.232)	(0.189)	(0.217)	(0.027)	(0.052)
Observations	93,026		93,026		93,026	

Notes: The table reports 2SLS estimates from the preferred specification of column (4) of Table 3, estimated on the common estimation sample of column (5). The automation and augmentation exposure measures are inverse-hyperbolic-sine transformed and standardized within this sample. The first row reports the point estimates, which are held constant across inference procedures. The subsequent rows report standard errors clustered by commuting zone, Conley spatial-HAC standard errors using Bartlett kernels at the indicated distance cutoffs, and standard errors two-way clustered by commuting zone and detailed occupation. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), include commuting-zone-by-broad-occupation fixed effects, and are weighted by 2005–07 FTE employment shares.

G Robustness to Major AI Innovation Hubs

AI innovation in the United States is geographically concentrated. Figure G.1 illustrates this concentration for the two patent populations underlying our empirical design: historical AI breakthrough patents granted between 1980 and 2000 in Panel A and contemporary AI patents granted between 2005 and 2021 in Panel B. Historical AI breakthroughs were concentrated particularly in Silicon Valley and along the Northeast Corridor, including New York and Boston, and many of these locations remain among the largest centers of contemporary AI patenting, as further documented in Appendix Tables G.2 and G.3. In each period, the three largest hubs alone account for more than 40 percent of the corresponding patents in our geographically assigned sample.

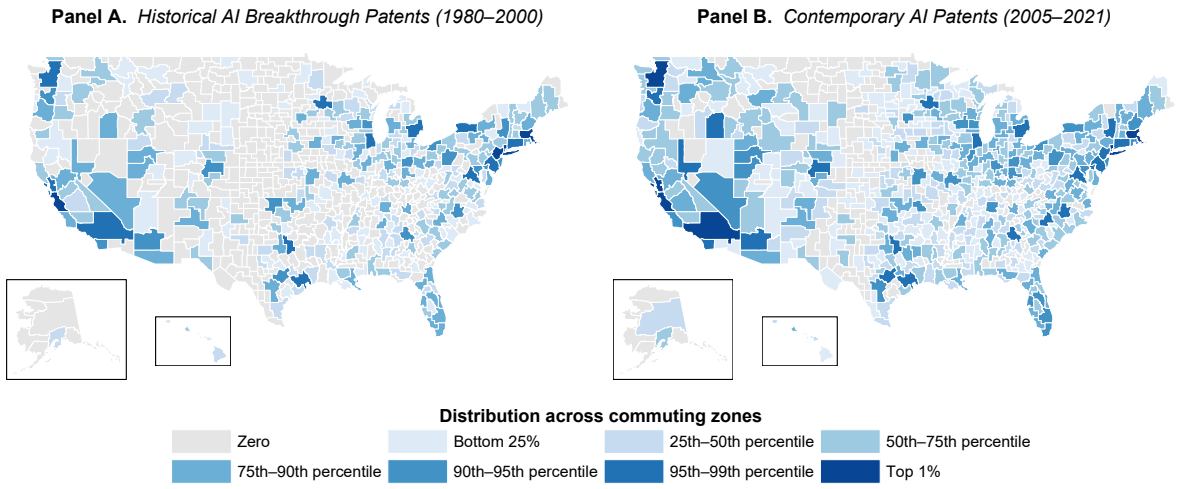


Figure G.1: Geographic Concentration of Historical and Contemporary AI Innovation

Notes: Panel A maps AI breakthrough patents granted during 1980–2000, and Panel B maps AI patents granted during 2005–2021. Patent origins are assigned to commuting zones using the first listed assignee location when a valid U.S. commuting-zone assignment is available and otherwise the first inventor location. Patents classified as foreign or lacking a valid U.S. commuting-zone assignment are excluded. Gray denotes commuting zones with no relevant patents. Among commuting zones with positive counts, darker shading indicates a higher percentile rank within the corresponding period-specific distribution.

This concentration raises two distinct concerns. First, the identifying variation in the instruments may depend disproportionately on breakthroughs originating in a small number of historically dominant AI hubs. Second, the estimated labor-market effects may be disproportionately influenced by contemporary AI hubs, where exposure is unusually high and local economic conditions may differ systematically from those elsewhere. Table G.1 examines these two concerns separately.

Columns (2)–(3) address the first concern by reconstructing the historical instruments after excluding the three and ten largest historical AI breakthrough hubs, respectively, from the set of donor commuting zones. The estimation sample and contemporary exposure measures remain

Table G.1: Robustness to Major Historical and Contemporary AI Hubs

		IV donor exclusions		Sample exclusions	
	Baseline	Top 3	Top 10	Top 5	Top 30
Panel A: Wage bill change					
AI Automation	-0.248***	-0.281***	-0.306***	-0.238***	-0.245***
	(0.019)	(0.019)	(0.022)	(0.018)	(0.023)
AI Augmentation	0.145***	0.168***	0.185***	0.118***	0.103***
	(0.026)	(0.026)	(0.029)	(0.024)	(0.026)
Panel B: Employment change					
AI Automation	-0.231***	-0.260***	-0.287***	-0.222***	-0.230***
	(0.016)	(0.017)	(0.020)	(0.016)	(0.022)
AI Augmentation	0.067***	0.089***	0.106***	0.050**	0.049**
	(0.020)	(0.020)	(0.022)	(0.020)	(0.023)
Panel C: Hourly wage change					
AI Automation	-0.011	-0.017**	-0.018**	-0.008	0.004
	(0.007)	(0.007)	(0.007)	(0.007)	(0.008)
AI Augmentation	0.079***	0.084***	0.091***	0.068***	0.043***
	(0.010)	(0.010)	(0.011)	(0.009)	(0.008)
KP F-statistic	711.01	817.14	971.39	592.67	533.94
Observations	93,026	93,026	93,026	91,216	83,137

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells based on the common estimation sample of column (5) of Table 3, with additional sample exclusions in columns (4)–(5). The dependent variables are log changes in wage bills and FTE employment and changes in average log hourly wages between 2005–07 and 2022–24. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within the baseline common estimation sample, with this standardization held fixed across all five columns. Column (1) reproduces column (4) of Table 3. Columns (2)–(3) reconstruct the historical instruments after excluding the three and ten largest historical AI breakthrough hubs, respectively, from the donor set. These hubs are ranked by the number of AI breakthrough patents granted during 1980–2000; the contemporary exposure measures and estimation sample remain unchanged. Columns (4)–(5) retain the baseline exposure measures and instruments but exclude from the estimation sample the five and thirty largest contemporary AI patent hubs, respectively, ranked by the number of unique AI patents granted during 2005–2021. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), include commuting-zone-by-broad-occupation fixed effects, and are weighted by 2005–07 FTE employment shares. Standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

unchanged. The resulting estimates preserve the baseline pattern and are somewhat larger in magnitude, while the instruments remain strong.

Columns (4)–(5) address the second concern by retaining the baseline exposure measures and instruments but excluding the five and thirty largest contemporary AI patent hubs, respectively, from the estimation sample. The automation employment and wage-bill estimates remain particularly stable, while the automation hourly-wage estimates remain small. The augmentation estimates become somewhat smaller as the largest contemporary hubs are removed but remain positive and statistically significant across all three outcomes. These results indicate that our main findings are not driven by a small number of dominant AI innovation hubs.

Table G.2: Major Historical AI Breakthrough Hubs, 1980–2000

(Commuting zones excluded from the donor set in the major-hub robustness analysis)

	Commuting Zone	AI Breakthroughs	Share (%)	All Breakthroughs	AI Share (%)
1	New York, NY (19400)	5,062	17.6	23,941	21.1
2	San Jose, CA (37500)	4,486	15.6	22,047	20.3
3	Boston, MA (20500)	2,091	7.3	10,713	19.5
4	Newark, NJ (19600)	1,612	5.6	10,681	15.1
5	Los Angeles, CA (38300)	1,224	4.3	9,130	13.4
6	Hartford, CT (20901)	1,148	4.0	6,813	16.9
7	Washington, DC (11304)	1,091	3.8	7,168	15.2
8	Chicago, IL (24300)	1,074	3.7	10,158	10.6
9	San Francisco, CA (37800)	1,005	3.5	6,326	15.9
10	Seattle, WA (39400)	991	3.4	3,342	29.7
	Top 10 total	19,784	68.8	110,319	17.9

Notes: Commuting zones are ranked by the number of AI breakthrough patents granted during 1980–2000, using the same geographic assignment as in the baseline instrument construction. For each commuting zone, the table reports the number of AI breakthrough patents, its share of all 28,763 geographically assigned U.S. AI breakthrough patents, the total number of breakthrough patents, and the share of those patents classified as AI-related. Of the 741 commuting zones, 311 contain at least one AI breakthrough patent. In the robustness analysis in columns (2)–(3) of Table G.1, the top three and top ten commuting zones are excluded from the set of donor commuting zones used to construct the alternative instruments; the recipient commuting zones, local exposure measures, and estimation sample remain unchanged. City and state labels correspond to the principal city within each commuting-zone polygon; CZ identifiers are reported in parentheses.

Table G.3: Major Contemporary AI Innovation Hubs, 2005–2021

(Commuting zones excluded from the estimation sample in the major-hub robustness analysis)

	Commuting Zone	AI Patents	Share (%)	All Patents	AI Share (%)
1	San Jose, CA (37500)	93,918	20.0	288,680	32.5
2	New York, NY (19400)	70,349	15.0	166,913	42.1
3	Seattle, WA (39400)	39,108	8.3	76,133	51.4
4	San Francisco, CA (37800)	38,847	8.3	111,766	34.8
5	Boston, MA (20500)	19,953	4.2	90,762	22.0
6	Los Angeles, CA (38300)	16,983	3.6	106,338	16.0
7	Washington, DC (11304)	11,988	2.5	45,815	26.2
8	Houston, TX (32000)	11,400	2.4	78,464	14.5
9	Newark, NJ (19600)	10,925	2.3	68,396	16.0
10	Chicago, IL (24300)	10,867	2.3	72,008	15.1
11	Atlanta, GA (09100)	10,822	2.3	40,288	26.9
12	San Diego, CA (38000)	9,796	2.1	65,512	15.0
13	Detroit, MI (11600)	7,962	1.7	68,625	11.6
14	Dallas, TX (33100)	6,932	1.5	41,351	16.8
15	Hartford, CT (20901)	6,614	1.4	44,434	14.9
16	Austin, TX (31201)	6,614	1.4	32,045	20.6
17	Albany, NY (18600)	4,599	1.0	29,905	15.4
18	Raleigh, NC (01701)	4,190	0.9	17,109	24.5
19	Reno, NV (37604)	4,118	0.9	11,721	35.1
20	Philadelphia, PA (19700)	4,087	0.9	27,380	14.9
21	Denver, CO (28900)	3,792	0.8	21,441	17.7
22	Charlotte, NC (00900)	3,674	0.8	13,624	27.0
23	Minneapolis, MN (21501)	3,589	0.8	37,235	9.6
24	Wilmington, DE (19800)	3,579	0.8	29,593	12.1
25	Portland, OR (38801)	3,373	0.7	18,572	18.2
26	Boise, ID (35801)	2,478	0.5	21,230	11.7
27	Phoenix, AZ (35001)	2,383	0.5	18,667	12.8
28	Cleveland, OH (15200)	2,360	0.5	21,905	10.8
29	Kansas City, MO (29502)	2,329	0.5	11,329	20.6
30	Rochester, NY (18000)	2,028	0.4	12,773	15.9
	Top 30 total	419,657	89.2	1,690,014	24.8

Notes: Commuting zones are ranked by the number of AI patents granted during 2005–2021, using the same geographic assignment as in the construction of the contemporaneous exposure measures. For each commuting zone, the table reports the number of AI patents, its share of all 470,498 geographically assigned U.S. AI patents, the total number of patents, and the share of those patents classified as AI-related. Of the 741 commuting zones, 534 contain at least one AI patent. In the robustness analysis in columns (4)–(5) of Table G.1, the top five and top thirty commuting zones are excluded from the estimation sample; the local exposure measures and instruments remain unchanged. City and state labels correspond to the principal city within each commuting-zone polygon; CZ identifiers are reported in parentheses.

H Robustness to Alternative Instrument Constructions

Earlier historical breakthrough window. Table H.1 addresses the concern that the baseline breakthrough window, which extends through 2000, may provide insufficient temporal separation from the contemporary exposure period beginning in 2005. We therefore reconstruct the instruments using only AI breakthroughs granted during 1980–1990. We treat this earlier window as an alternative rather than our preferred specification because it draws on a substantially smaller set of breakthroughs. The 2SLS estimates remain qualitatively consistent with the baseline results but are larger in magnitude. This is particularly pronounced for AI augmentation, whose employment estimate rises from 0.067 to 0.260, while the hourly-wage estimate changes much less. AI automation likewise produces larger negative employment and wage-bill estimates. Thus, while the magnitude of the estimates is noticeably sensitive to the historical window, their direction across both channels and all three outcomes is preserved.

Larger geographic exclusion buffer. The baseline instruments exclude historical breakthroughs originating in the focal commuting zone and the first two surrounding contiguity rings. Table H.2 adopts a more conservative exclusion rule by constructing the instruments only from breakthroughs located at least six contiguity rings away from the focal commuting zone. The resulting estimates are very similar to the baseline estimates in both sign and magnitude. The stability of the estimates under this wider buffer reduces the concern that the baseline results are driven by spatial spillovers or other economic linkages associated with geographically proximate historical innovation.

Intermediate-distance breakthrough exposure. As a complementary exercise, we construct the instruments exclusively from historical breakthroughs originating in rings 3–5, with the estimates reported in Table H.3. Whereas the preceding exercise deliberately excludes these rings, this specification instead omits more distant breakthroughs and uses only the nearest donor rings eligible under the baseline exclusion rule. The estimates preserve the baseline pattern and are broadly comparable in magnitude, although the augmentation employment estimate becomes somewhat smaller and less precise.

Alternative distance weighting. Table H.4 examines whether the results depend on the particular spatial-decay assumption used to construct the historical instruments. Instead of the baseline inverse-distance weights, $1/d$, the instruments weight breakthroughs using inverse-square distance weights, $1/d^2$, placing substantially greater weight on nearer donor locations outside the exclusion buffer. Complementing the preceding spatial exercises, the steeper distance weighting leaves the estimates nearly unchanged.

Relative AI exposure measures. Finally, Table H.5 repeats these instrument-construction exercises for the relative AI exposure measures using the preferred specification throughout. The

contemporary relative AI exposure measures and the corresponding total automation and augmentation exposure controls are held fixed across columns for comparison. For automation, the negative employment and wage-bill estimates are preserved across all instrument constructions, while the hourly-wage estimates remain comparatively small. For augmentation, the positive hourly-wage estimates are similarly stable, while the employment estimate is more sensitive, especially to the earlier historical breakthrough window, mirroring the sensitivity of the level augmentation estimates in Table H.1. This is also reflected in the corresponding augmentation wage-bill estimates, although they remain nonnegative across all constructions despite the noticeable changes in magnitude.

Table H.1: Robustness to an Earlier Historical AI Breakthrough Window

	Baseline IV 1980–2000	Earlier-window IV 1980–1990
Panel A: Wage bill change		
AI Automation	-0.248***	-0.514***
	(0.019)	(0.030)
AI Augmentation	0.145***	0.327***
	(0.026)	(0.037)
Panel B: Employment change		
AI Automation	-0.231***	-0.485***
	(0.016)	(0.025)
AI Augmentation	0.067***	0.260***
	(0.020)	(0.029)
Panel C: Hourly wage change		
AI Automation	-0.011	-0.030***
	(0.007)	(0.010)
AI Augmentation	0.079***	0.092***
	(0.010)	(0.012)
KP F-statistic	711.01	613.96
Observations	93,026	93,026

Notes: The table reports 2SLS estimates comparing the preferred specification under the baseline and earlier historical breakthrough windows. The first column uses the baseline historical AI breakthrough instruments constructed from patents granted during 1980–2000. The second column replaces these instruments with otherwise analogous instruments constructed exclusively from AI breakthrough patents granted during 1980–1990. Contemporary local AI automation and augmentation exposure continue to cover 2005–2021 and are unchanged across columns. Both instrument sets use the baseline 90th-percentile patent–task and patent–tool similarity threshold and the same two-ring exclusion buffer, with donor commuting zones from ring 3 onward. Both columns and all panels use the common estimation sample of column (5) of Table 3. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this common sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing), include commuting-zone-by-broad-occupation fixed effects, and are weighted by 2005–07 FTE employment shares. Standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table H.2: Robustness to a Five-Ring Exclusion Buffer

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.152***	-0.149***	-0.253***	-0.262***	-0.110***
	(0.034)	(0.029)	(0.022)	(0.019)	(0.022)
AI Augmentation	0.167***	0.215***	0.171***	0.160***	0.105***
	(0.051)	(0.047)	(0.027)	(0.027)	(0.024)
Panel B: Employment change					
AI Automation	-0.147***	-0.131***	-0.233***	-0.242***	-0.142***
	(0.029)	(0.025)	(0.019)	(0.016)	(0.018)
AI Augmentation	0.138***	0.137***	0.092***	0.080***	0.042**
	(0.041)	(0.039)	(0.021)	(0.020)	(0.019)
Panel C: Hourly wage change					
AI Automation	-0.014*	-0.021***	-0.017**	-0.015**	0.010
	(0.008)	(0.007)	(0.008)	(0.008)	(0.009)
AI Augmentation	0.034**	0.071***	0.085***	0.084***	0.075***
	(0.013)	(0.013)	(0.011)	(0.011)	(0.011)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	423.97	475.93	417.25	570.92	545.95
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using historical AI breakthrough instruments constructed with a five-ring exclusion buffer. The excluded instruments use AI breakthrough patents granted during 1980–2000 originating at least six contiguity rings from the focal commuting zone. Contemporary local AI automation and augmentation exposure remain unchanged. All columns and panels use the common estimation sample of column (5) of Table 3. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this common sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table H.3: Robustness to Breakthrough Exposure from Rings 3–5

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.120***	-0.135***	-0.239***	-0.228***	-0.088***
	(0.039)	(0.030)	(0.026)	(0.024)	(0.026)
AI Augmentation	0.059	0.181***	0.121***	0.118***	0.074***
	(0.069)	(0.049)	(0.032)	(0.028)	(0.026)
Panel B: Employment change					
AI Automation	-0.121***	-0.120***	-0.223***	-0.213***	-0.122***
	(0.034)	(0.028)	(0.025)	(0.023)	(0.025)
AI Augmentation	0.031	0.106***	0.045*	0.042*	0.012
	(0.055)	(0.039)	(0.026)	(0.023)	(0.022)
Panel C: Hourly wage change					
AI Automation	-0.004	-0.016**	-0.011	-0.009	0.015*
	(0.010)	(0.008)	(0.008)	(0.008)	(0.008)
AI Augmentation	0.029	0.067***	0.082***	0.078***	0.070***
	(0.021)	(0.014)	(0.012)	(0.010)	(0.010)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	55.66	343.68	326.90	393.15	334.20
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using historical AI breakthrough instruments constructed exclusively from contiguity rings 3–5 around the focal commuting zone. The focal commuting zone and rings 1–2 remain excluded, while breakthrough exposure from ring 6 onward is omitted. Contemporary local AI automation and augmentation exposure remain unchanged. All columns and panels use the common estimation sample of column (5) of Table 3. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this common sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table H.4: Robustness to Inverse-Square Distance Weighting

	(1)	(2)	(3)	(4)	(5)
Panel A: Wage bill change					
AI Automation	-0.144***	-0.146***	-0.244***	-0.237***	-0.090***
	(0.034)	(0.028)	(0.021)	(0.020)	(0.022)
AI Augmentation	0.148***	0.210***	0.154***	0.132***	0.081***
	(0.048)	(0.046)	(0.027)	(0.026)	(0.024)
Panel B: Employment change					
AI Automation	-0.144***	-0.131***	-0.228***	-0.222***	-0.125***
	(0.030)	(0.025)	(0.020)	(0.018)	(0.020)
AI Augmentation	0.114***	0.136***	0.077***	0.056***	0.021
	(0.040)	(0.039)	(0.022)	(0.021)	(0.019)
Panel C: Hourly wage change					
AI Automation	-0.006	-0.016**	-0.010	-0.009	0.016**
	(0.008)	(0.007)	(0.008)	(0.007)	(0.008)
AI Augmentation	0.035**	0.066***	0.081***	0.078***	0.070***
	(0.014)	(0.012)	(0.010)	(0.010)	(0.010)
State FE	✓				
CZ FE		✓	✓		
Broad occupation FE			✓		
CZ × broad occupation FE				✓	✓
College share					✓
Routine intensity					✓
KP F-statistic	120.63	699.53	617.67	752.31	764.46
Observations	93,026	93,026	93,026	93,026	93,026

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using historical AI breakthrough instruments that weight exposure by the inverse square of contiguity-ring distance, $1/d^2$, rather than the baseline inverse-distance weight, $1/d$. As in the baseline construction, the focal commuting zone and rings 1–2 are excluded, and historical breakthrough exposure from ring 3 onward is used. Contemporary local AI automation and augmentation exposure remain unchanged. All columns and panels use the common estimation sample of column (5) of Table 3. AI automation and augmentation exposure are inverse-hyperbolic-sine transformed and standardized within this common sample. All specifications control for non-AI automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing). Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

Table H.5: Robustness of Relative AI Effects to Alternative Instrument Constructions

	Baseline	Earlier window	Five-ring buffer	Rings 3–5	Inverse-square weights
IV window	1980–2000	1980–1990	1980–2000	1980–2000	1980–2000
Spatial construction	$d \geq 3,\ 1/d$	$d \geq 3,\ 1/d$	$d \geq 6,\ 1/d$	$3 \leq d \leq 5,\ 1/d$	$d \geq 3,\ 1/d^2$
Panel A: Wage bill change					
Relative AI Automation	-0.030***	-0.089***	-0.025***	-0.044***	-0.036***
	(0.008)	(0.008)	(0.007)	(0.011)	(0.010)
Relative AI Augmentation	0.009	0.061***	0.002	0.022**	0.015*
	(0.007)	(0.008)	(0.007)	(0.009)	(0.008)
Panel B: Employment change					
Relative AI Automation	-0.026***	-0.082***	-0.021***	-0.040***	-0.032***
	(0.008)	(0.007)	(0.007)	(0.011)	(0.010)
Relative AI Augmentation	-0.027***	0.023***	-0.034***	-0.012	-0.021**
	(0.007)	(0.006)	(0.007)	(0.010)	(0.009)
Panel C: Hourly wage change					
Relative AI Automation	0.005*	-0.003	0.006**	0.004	0.004
	(0.003)	(0.003)	(0.003)	(0.003)	(0.003)
Relative AI Augmentation	0.030***	0.038***	0.029***	0.031***	0.031***
	(0.004)	(0.003)	(0.004)	(0.003)	(0.003)
KP F-statistic	464.06	669.83	377.25	337.83	498.86
Observations	88,539	88,539	88,539	88,539	88,539

Notes: The table reports 2SLS estimates for occupation–commuting-zone cells using alternative historical AI breakthrough instrument constructions for the relative AI exposure measures. Relative AI automation and augmentation exposure are defined as the AI shares of total exposure within each channel and instrumented using the corresponding AI shares of total historical-breakthrough exposure. The contemporary relative AI exposure measures are held fixed across columns; only the corresponding historical instrument constructions vary. Column (1) uses AI breakthroughs granted during 1980–2000, excludes the focal commuting zone and rings 1–2, and weights exposure from ring 3 onward by $1/d$. Column (2) uses the earlier 1980–1990 breakthrough window with the same spatial construction. Column (3) expands the exclusion buffer through ring 5, column (4) uses only breakthrough exposure from rings 3–5, and column (5) replaces the baseline $1/d$ weights with $1/d^2$. All columns use a common estimation sample requiring every alternative relative instrument to be defined. Relative AI automation and augmentation are standardized within this common sample. All specifications control for total automation and augmentation exposure and baseline workforce characteristics (shares below age 30, above age 55, female, foreign-born, and employed in manufacturing) and include commuting-zone-by-broad-occupation fixed effects. Estimates are weighted by 2005–07 FTE employment shares, and standard errors clustered by commuting zone are reported in parentheses. The Kleibergen–Paap F-statistic is the cluster-robust weak-identification statistic for the excluded instruments. ***, **, and * indicate significance at the 1, 5, and 10 percent levels.

I Details and Derivations for the Conceptual Framework

This appendix provides additional details and derivations for the conceptual framework in Section 2. We first derive the task-level unit costs and then distinguish explicitly between the technological feasibility of autonomous production and its cost-minimizing use. We finally derive the conditional labor-demand and physical-productivity effects underlying the labor-market implications in the main text. These derivations do not impose unconditional signs on equilibrium employment, wages, or the wage bill.

I.1 Unit Costs

Recall from equation (2) that task i may be produced through the labor–tool method or, when technologically feasible, through autonomous production. Output under the two methods is

$$y_{oc}^{L}(i) = \gamma_{o}^{L}(i)\ell_{oc}(i)^{1-\kappa}\left\{\Gamma_{oc}x_{oc}(i)\right\}^{\kappa}, \qquad y_{oc}^{K}(i) = A_{oc}^{K}\gamma_{o}^{K}(i)k_{oc}(i), \tag{30}$$

where $0 < \kappa < 1$. The labor–tool method is technologically available throughout the task space. Let $\bar{I}_o \in [0,1]$ denote the occupation-specific technological feasibility frontier for autonomous production. Autonomous production is technologically feasible for tasks $i \leq \bar{I}_o$, whereas tasks with $i > \bar{I}_o$ currently require labor. We take this technological feasibility frontier to be common across commuting zones because it captures the tasks that existing technology is capable of performing autonomously rather than the extent of its local implementation.

Let w_{oc} denote the occupational wage, R_{oc}^{X} the unit cost of the tool input, and R_{oc}^{K} the rental cost of autonomous capital. Consider first the labor–tool method. Its unit cost is obtained from

$$q_{oc}^{L}(i) = \min_{\ell_{oc}(i),x_{oc}(i)} \left\{w_{oc}\ell_{oc}(i) + R_{oc}^{X}x_{oc}(i)\right\}, \tag{31}$$

subject to

$$\gamma_{o}^{L}(i)\ell_{oc}(i)^{1-\kappa}\left\{\Gamma_{oc}x_{oc}(i)\right\}^{\kappa} \geq 1. \tag{32}$$

Because the production function is Cobb–Douglas and homogeneous of degree one in labor and the tool input, cost minimization implies

$$w_{oc}\ell_{oc}(i) = (1-\kappa)q_{oc}^{L}(i), \qquad R_{oc}^{X}x_{oc}(i) = \kappa q_{oc}^{L}(i). \tag{33}$$

Hence,

$$\ell_{oc}(i) = \frac{(1-\kappa)q_{oc}^{L}(i)}{w_{oc}}, \qquad x_{oc}(i) = \frac{\kappa q_{oc}^{L}(i)}{R_{oc}^{X}}. \tag{34}$$

Substituting equation (34) into the unit-output condition (32) gives

$$1 = \gamma_{o}^{L}(i)\left[\frac{(1-\kappa)q_{oc}^{L}(i)}{w_{oc}}\right]^{1-\kappa}\left[\Gamma_{oc}\frac{\kappa q_{oc}^{L}(i)}{R_{oc}^{X}}\right]^{\kappa}, \tag{35}$$

and since the exponents sum to one, solving for $q_{oc}^L(i)$ yields

$$q_{oc}^L(i) = \frac{1}{\gamma_o^L(i)} \left(\frac{w_{oc}}{1-\kappa} \right)^{1-\kappa} \left(\frac{R_{oc}^X}{\kappa \Gamma_{oc}} \right)^{\kappa}. \tag{36}$$

Conditional on autonomous production being technologically feasible, producing one unit of task output requires

$$k_{oc}(i) = \frac{1}{A_{oc}^K \gamma_o^K(i)}. \tag{37}$$

The corresponding conditional unit cost of autonomous production is therefore

$$q_{oc}^K(i) = \frac{R_{oc}^K}{A_{oc}^K \gamma_o^K(i)}. \tag{38}$$

For tasks above the technological feasibility frontier, $i > \bar{I}_o$, autonomous production is not technologically feasible. Equivalently, its effective unit cost can be treated as infinite. The minimum cost of producing task i can therefore be written as

$$q_{oc}(i) = \begin{cases} \min \left\{ q_{oc}^L(i), q_{oc}^K(i) \right\}, & i \leq \bar{I}_o, \\ q_{oc}^L(i), & i > \bar{I}_o. \end{cases} \tag{39}$$

Equation (39) separates technological feasibility from cost-minimizing task assignment. For $i \leq \bar{I}_o$, both production methods are technologically feasible and the method actually used is determined by their relative unit costs.

Equations (36) and (38) give the unit costs reported in equation (4) in the main text. Conditional on the corresponding production method being technologically available and used,

$$\frac{\partial \ln q_{oc}^L(i)}{\partial \ln \Gamma_{oc}} = -\kappa, \qquad \frac{\partial \ln q_{oc}^L(i)}{\partial \ln R_{oc}^X} = \kappa, \qquad \frac{\partial \ln q_{oc}^K(i)}{\partial \ln A_{oc}^K} = -1. \tag{40}$$

Thus, greater tool productivity or a lower unit cost of the tool input reduces the cost of producing a task using labor and tools, while greater productivity of autonomous technology lowers the conditional cost of autonomous production.

I.2 Technological Feasibility and Cost-Minimizing Task Assignment

Within the technologically feasible set, $i \leq \bar{I}_o$, task i is performed through autonomous production whenever

$$q_{oc}^K(i) \leq q_{oc}^L(i), \tag{41}$$

or, using equations (36) and (38),

$$\frac{\gamma_o^L(i)}{\gamma_o^K(i)} \leq \frac{A_{oc}^K}{R_{oc}^K} \left(\frac{w_{oc}}{1-\kappa} \right)^{1-\kappa} \left(\frac{R_{oc}^X}{\kappa \Gamma_{oc}} \right)^{\kappa}. \tag{42}$$

Define relative labor productivity across tasks as

$$\Psi_o(i) \equiv \frac{\gamma_o^L(i)}{\gamma_o^K(i)}. \tag{43}$$

As in the standard task framework, suppose that labor's comparative advantage increases with the task index,

$$\Psi_o'(i) > 0. \tag{44}$$

This monotonicity implies that the relative-cost comparison can be summarized by an economic assignment threshold $\widetilde{I}_{oc}$. When the threshold is interior, it satisfies

$$q_{oc}^L(\widetilde{I}_{oc}) = q_{oc}^K(\widetilde{I}_{oc}), \tag{45}$$

or equivalently,

$$\Psi_o(\widetilde{I}_{oc}) = \frac{A_{oc}^K}{R_{oc}^K}\left(\frac{w_{oc}}{1-\kappa}\right)^{1-\kappa}\left(\frac{R_{oc}^X}{\kappa\Gamma_{oc}}\right)^{\kappa}. \tag{46}$$

The two thresholds therefore have different meanings. The technological feasibility frontier $\bar{I}_o$ determines which tasks autonomous technology is currently capable of performing. The economic assignment threshold $\widetilde{I}_{oc}$ determines, conditional on technological feasibility, the range of tasks for which autonomous production would be less costly than the labor–tool method. Actual task assignment is limited by both conditions, so the effective automation frontier used in Section 2 is

$$I_{oc} = \min\left\{\bar{I}_o, \widetilde{I}_{oc}\right\}. \tag{47}$$

Tasks with $i < I_{oc}$ are performed autonomously, while tasks with $i > I_{oc}$ are performed by workers using complementary tools.[20]

Changes in the effective automation frontier can therefore arise either from an expansion in the set of tasks that autonomous technology is capable of performing or from improvements in the productivity or reductions in the cost of autonomous production that make it the cost-minimizing method for additional tasks. Improvements in autonomous technology that affect only tasks already performed autonomously, without changing I_{oc}, instead represent a deepening of automation, a margin we do not consider further here.

Because $q_{oc}^L(i)$ enters the relative-cost condition in equation (46), improvements in complementary tools may in principle also lower the economic assignment threshold $\widetilde{I}_{oc}$ by making production using labor and tools relatively less costly. The main-text comparative statics abstract from this induced task-reassignment margin and isolate the direct tool-complementarity effect conditional on the task continuing to be performed by labor.

[20] For simplicity, the expressions above describe an interior economic threshold. Corner cases can be represented by truncating $\widetilde{I}_{oc}$ at the boundaries of the task space.

I.3 Conditional Labor Demand and Productivity

The labor-demand expressions in Section 2 follow directly from the conditional factor demands of the labor–tool production method. To produce $y_{oc}(i)$ units of a labor-performed task, total expenditure is $q_{oc}^{L}(i)y_{oc}(i)$. Because labor accounts for the share $1-\kappa$ of expenditure under the Cobb–Douglas technology, equation (33) implies

$$w_{oc}\ell_{oc}(i) = (1-\kappa)q_{oc}^{L}(i)y_{oc}(i). \tag{48}$$

Equation (48) gives equation (6) in the main text. Aggregating over labor-performed tasks yields equation (7). Holding task output and factor prices fixed, applying Leibniz's rule to its lower integration bound gives the direct automation effect in equation (8).

For tool augmentation, equations (36) and (6) additionally imply, conditional on task output, wages, and task assignment,

$$\frac{\partial \ln \ell_{oc}(i)}{\partial \ln \Gamma_{oc}} = -\kappa, \qquad \frac{\partial \ln \ell_{oc}(i)}{\partial \ln R_{oc}^{X}} = \kappa, \qquad i > I_{oc}. \tag{49}$$

Thus, greater tool productivity or a lower unit cost of the tool input reduces the amount of labor demanded to produce a given quantity of a task that remains performed by labor.

A higher Γ_{oc} simultaneously raises the physical marginal product of labor. From equation (30),

$$MPL_{oc}^{L}(i) \equiv \frac{\partial y_{oc}^{L}(i)}{\partial \ell_{oc}(i)} = (1-\kappa)\gamma_{o}^{L}(i)\ell_{oc}(i)^{-\kappa}\left\{\Gamma_{oc}x_{oc}(i)\right\}^{\kappa}. \tag{50}$$

Holding labor and the tool input fixed therefore gives

$$\frac{\partial \ln MPL_{oc}^{L}(i)}{\partial \ln \Gamma_{oc}} = \kappa > 0, \qquad i > I_{oc}. \tag{51}$$

Tool augmentation can therefore raise workers' physical productivity while reducing the amount of labor demanded for a given level of task output. These are conditional technological effects rather than equilibrium employment or wage effects. Lower production costs may increase task and occupational output, while equilibrium wages depend on the value rather than the physical marginal product of labor and therefore also reflect adjustments in task prices and other equilibrium responses.