

ARTICLE

Augmented, not Automated: AI Impact on Jobs in Europe

By: Charles Zhang & Dennis Schoenmaker



Executive Summary

01

Creative destruction, again.

Like every prior technology wave, AI creates new roles and new ways of working as it retires routine ones. AI-skill job postings have climbed in every major economy since ChatGPT was introduced, and job demand is tilting from routine work toward judgment-heavy roles that AI is unlikely to replace.

58%

Share of S&P 500 Earnings Calls Mentioning AI (Q4 2025)

02

Complement, not substitute.

Because non-routine judgment is what AI cannot yet do, it augments these workers rather than replaces them. Both corporate signalling and workflow data confirm this: AI came up in a workforce context on 58% of S&P 500 earnings calls in Q4 2025, with complement-style language outweighing substitute-style by roughly eight to one (Citadel Securities, 2026). Across 100,000+ GitHub developers, AI-generated initial output surged around seventeen fold, while final product releases grew only about 30% because more human judgment is needed at every step between initial output and final release (Demirer et al., 2026).

17x

Growth in AI-Generated Developer Output

03

Exposure alone doesn't predict harm, adaptability does.

AI exposure varies widely across European cities, but exposure and adaptability tend to move together, so the most exposed markets are often the best equipped to absorb the disruption. In the most adaptable markets – the UK and France among them – causal evidence shows company formation in AI-exposed industries accelerating after ChatGPT and strengthening over time.

04

The signal is already visible in real estate.

In London – one of Europe's leading AI markets – AI-exposed occupiers are taking a rising and increasingly dispersed share of new office take-up, with demand diffusing outward from the original cluster core.

Every Technology Wave Creates New Ways of Working and New Jobs


Every wave of transformative technology destroys some jobs, reshapes many others, and creates roles that nobody could have foreseen. AI is no exception, and history offers a template for reading what is happening today.

Tracking job titles through Google Ngram data (Figure 1) across four technology eras shows a consistent rhythm: Each title rose, peaked and faded as the underlying work changed, replaced by a new title describing work the previous era did not need. Obsolescence has consistently been matched by new gains, not simply losses.


The same pattern holds within surviving job titles. Job ads for financial managers (Figure 2) show the task mix shifting steadily from routine work toward non-routine judgment throughout previous tech waves. AI is likely to absorb the remaining routine work even further but unlikely to erode non-routine judgment given its current ability.

¹ Google Ngram tracks how frequently a term appears per 1 million words across all published books and texts in the US, making it a widely used tool for tracking how job titles evolve over time (see Autor et al., 2022). The higher the score, the more such job title is being mentioned in a particular time.

Figure 1. Prominence of Job Titles Measured by Google Ngram



Source: CBRE Research, Google Ngram

Figure 2. Required Tasks in Job Ads for Financial Managers ²

Source: CBRE Research, Atalay et al. (2020)

² This figure is reconstructed from the dataset assembled in Atalay (2020), which analyses tasks over 7.8 million job ads published in the US from 1950 to 2000.


The same signals are visible today, faster than before. Since ChatGPT's release, AI-skill job postings, based on data from Indeed, have climbed steeply across every major economy we track (Figure 3) – some genuinely new roles, some AI fluency bolted onto existing ones. Either way, AI creates new AI-native labour demand, as in every prior wave.

Demand composition is also shifting in a familiar direction: UK postings for less-routine roles have outpaced the average market rate since ChatGPT's release (Figure 4), holding steady even in the most 'exposed' occupation, such as software development. The same routine-biased shift seen in every wave is what is unfolding in European labour markets today.

³ United Kingdom, Germany, Italy, France, Netherlands, Ireland, Australia, United States.


⁴ Chart shows the spread between each sector's job-postings index and the UK all-sector average, rebased to zero in November 2022, the month ChatGPT launched.

Figure 3. Share of Job Postings Containing AI Terms by Country



Source: CBRE Research, Indeed Hiring Lab

Figure 4. UK Job Postings by Sector Relative to the Overall Market Since ChatGPT ⁴



Source: CBRE Research, Indeed Hiring Lab

The Evidence That AI Complements Workers

That shift in demand is the first clue to a bigger point: AI is set to complement workers, not replace them. Non-routine jobs' continued robust demand implies AI's limited ability to tackle unseen, judgment-heavy problems. Two independent data sources (S&P 500 earnings calls and GitHub developer log) tell the same story: Businesses are deploying AI to make existing staff more productive, not to replace them wholesale. One tracks corporate communication, the other developer behaviour.

Figure 5 shows that AI came up in a workforce context on 58% of S&P 500 earnings calls in the fourth quarter of 2025, up from just 17% in the first quarter of 2021, and when executives frame AI's workforce impact, complement-style language outweighs substitute-style language by roughly eight to one.

Workflow data tells the same story. Figure 6 tracks productivity gains from more than 100,000 GitHub developers across six stages of the coding pipeline for three generations of AI coding tool. The gain narrows at every stage: At the earliest stage, the volume of code generated has surged by roughly 17 times, but that gain does not carry through to the final stage, where the volume that ships as a finished release has grown by only about 30%. This is because even the most advanced models cannot reliably judge whether the code they write is useful or sound. AI therefore frees workers from routine coding tasks so they can focus on the more abstract judgment calls it still cannot make.


8:1

Complement vs. Substitute
Language Ratio

17x vs. 30%


GitHub Developer Pipeline:
Initial Code Output vs. Final
Shipped Releases

Figure 5. Executive Framing of the Impact of AI on the Workforce on Earnings Calls



Source: CBRE Research, Citadel Securities, LSEG StreetEvents

Figure 6. Productivity Gains of AI Coding Tools Across Production Stages ⁵



CBRE Research, Demirer et al. (2026)

⁵ Technical terms specific for coding production in the chart: lines changed counts the number of lines of code added or edited using AI; distinct files are unique files AI touched; commits are saved changes made by AI to the codebase; pull requests are proposed changes submitted for review; distinct repos counts how many separate codebases a developer's AI-assisted commits touched; and releases are the finished code that ships to production.

The Complement Effect Has Room Left to Run


If AI is a complement rather than a substitute, the next question is how much further that complementary effect can run. The cost and organisational evidence both point to significant room left to grow.

The cost of a fixed AI capability has fallen by roughly 1,000 times in three years due to model improvements (Figure 7). Cheaper tokens let AI absorb a larger share of the routine workload at a given budget, freeing human workers to concentrate on the non-routine judgment that actually drives output.

Cost alone does not deliver the productivity gain immediately. Historical technology adoption follows a J-curve (Figure 8): Productivity dips at the outset as organisations absorb the disruption, and the payoff only materialises once firms make complementary changes – reorganising workflows, retraining staff and rebuilding processes around the new tool. AI adoption appears to be following the same curve, and most organisations are still early in that reorganisation phase, which implies the larger productivity payoff may still be ahead of us. However, human input remains an integral part of that process.


⁶ This figure documents the evolution of token pricing over time. The lower the number, the less it costs to have a model complete the mentioned tasks.

Figure 7. Cheapest LLM Model Meeting Each Capability Bar Over Time ⁶



Source: CBRE Research, Epoch AI

Figure 8. Productivity J-Curve of AI Adoption ⁷



Source: CBRE Research, Brynjolfsson et al. (2019)

⁷ The horizontal axis represents time elapsed since a new technology is adopted; the vertical axis represents the relative level of productivity. The curve is stylised to illustrate this general pattern rather than measured AI-specific data, so the values on both axes are illustrative only.


AI Exposure and Adaptability Across Europe

Translating this from a labour-market story into a real estate story starts with each location's employment profile. We take a bottom-up approach that first estimates how many tasks within an occupation AI can comfortably replace, then maps that occupation's weight within each industry (for example, how many finance-industry roles are data scientists), and finally applies the local distribution of workers across industries.

This approach gives each location a theoretical AI exposure score – the share of job tasks AI is technically capable of performing (Figure 9). Comparing that theoretical exposure against actual AI usage – the share of those tasks AI is being used for today – reveals a persistent gap between potential and reality. In City of London, for example, AI could in principle perform roughly half of the tasks currently carried out in the job base, yet today it handles only around a fifth. That gap does not necessarily signal future workforce reduction; it is also a store of untapped capacity that has yet to be freed up for augmenting the existing workforce.

⁸ Theoretical AI exposure is the estimated share of job tasks in a location's workforce that AI is technically capable of performing; actual AI usage is the estimated share of those tasks for which AI is being used today.

Figure 9. Theoretical AI Exposure vs. Actual Usage in Major European Cities ⁸




Source: CBRE Research, Eloundou et al. (2022), Anthropic, BLS, Eurostat, ONS

More crucially, markets exposed to AI also tend to be adaptable to it. We build a market-level AI-adaptability score from factors that determine how easily a workforce can pivot away from exposed tasks, including measures of skill transferability and income (a proxy for the financial cushion and access to retraining that ease that transition). Layering that score onto the exposure measure (Figure 10) shows that exposure and adaptability tend to move together across European markets. Cities with the highest theoretical exposure to AI also tend to have the workforce characteristics needed to absorb that disruption constructively. High exposure, in other words, does not automatically translate into high vulnerability. In the markets best positioned to adapt, it looks instead like an opportunity to lead the next wave of AI-driven growth.

⁹ AI exposure and AI adaptive capacity values are log demeaned. A higher score on either axis means a market's exposure, or its adaptability, is further above the European average.

Figure 10. AI Exposure and AI Adaptability Across European Markets ⁹




Source: CBRE Research, Eloundou et al. (2022), Anthropic, BLS, Eurostat, ONS

AI Is Already Fuelling New Business Formation

If high-adaptability cities are positioned to benefit, we should expect to see that show up in real economic activity, and the business formation data bears this out. Exploiting cross-industry variation in AI exposure as a natural experiment, we test its causal effect on new business formation, benchmarking each industry's post-ChatGPT formation rate against its own pace immediately before the launch. In both the UK and France, two of Europe's more adaptable labour markets, company formation in AI-exposed industries accelerates sharply after ChatGPT's release, with the effect strengthening over time (Figures 11 and 12). Each point in Figure 11 and 12 tracks how much more, or less, AI-exposed industries are forming new firms than their own pace immediately before ChatGPT's launch; the shaded band is the 95% confidence interval around that estimate, and the dashed lines show the fitted slope before and after the release. The same pattern shows up in other AI frontrunners in Europe, such as the Netherlands. The mechanism is intuitive: Where AI can absorb routine work and drive down the cost of capability, the barriers to starting and scaling a business fall, so heavy AI exposure could show up as a wave of new firms rather than as job losses.


¹⁰ This figure compares business formation in industries with high versus low AI exposure, each month. The y-axis shows how much more (or less) new business formation an industry sees for each unit increase in its AI exposure, relative to that baseline, in percentage points; the shaded band is the 95% confidence interval. Dark and green lines show the average trend before and after ChatGPT's release, respectively.

Figure 11. Effect of AI Exposure on New Business Formation – UK ¹⁰



Source: CBRE Forecasting, Companies House

Figure 12. Effect of AI Exposure on New Business Formation – France



Source: CBRE Research, SIRENE

London, a Live Case Study in AI and Real Estate


New company formation and higher labour productivity both eventually need somewhere to work. London offers a live case study of how AI-driven economic activity may already be showing up in commercial real estate demand.

To study this phenomenon, we use CBRE's proprietary leasing data to give each occupier an AI-exposure score based on the granular sector in which their main business operates. We then overlay a 40-metre grid across central London and colour each tile by the AI-exposure of the leases signed nearby – the deeper the red, the higher the AI-exposure at that location. Comparing 2016–19 with 2023–26 shows a rising share of new office leasing going to more AI-exposed occupiers (Figures 13 and 14). We cannot fully separate whether this reflects newly formed AI-native firms taking space or established occupiers adapting and relocating, but either explanation points to the same underlying conclusion: AI-exposed occupiers are claiming a growing share of London's new office take-up.

40m

Grid Resolution Used to
Map AI-Exposure Across
Central London


Figure 13. AI-Exposure of New Office Lettings in London – 2016–19



Source: CBRE Research

The spatial pattern is also shifting. The core of the Knowledge Quarter cluster cools slightly between the two periods, while the districts ringing it heat up. AI-exposed demand that once concentrated inside the original cluster core appears to be diffusing outward into surrounding submarkets – the kind of spillover effect you would expect to see as a successful cluster outgrows its original footprint.

Figure 14. AI-Exposure of New Office Lettings in London – 2023–26



Source: CBRE Research

London is a useful case study precisely because it combines high theoretical AI exposure with high adaptability. The direction of travel – more leasing going to AI-exposed occupiers – is exactly what the earlier sections of this paper would predict: AI is augmenting workers, not displacing them. As it absorbs more routine tasks, it creates more room for non-routine work, keeping the existing workforce in place while leaving scope for further expansion ahead.

Looking Ahead

The evidence assembled here supports a single, consistent narrative. AI is following the same creative destruction arc as every prior technology wave: It creates new roles, new ways of working and new businesses.

It also disrupts routine ways of working. Importantly, however, the destruction is unlikely to happen to the wider labour force because AI is unlikely to affect the increasingly important non-routine tasks.

However, it is important to note that AI's impacts are routine-biased: AI absorbs routine tasks first, so to remain valuable, workers need to shift their effort toward the non-routine, judgment-heavy work that AI cannot yet reliably perform. The highly adaptable regions coincide with the highly exposed regions, and these are the places with the workforce and organisational flexibility to turn that exposure into new business formation, new roles, and new demand for space.

London's office market already shows early, tangible signs of this shift, with AI-exposed occupiers taking a growing and spatially diffusing share of new leasing. This is the starting point for a companion paper – the 'so what' for real estate – which carries the similar bottom-up approach into the office sector itself. Building up from a granular, firm-level dataset, it will trace how the evolution of employment, together with the pace of AI adoption and diffusion, feeds through into office demand and location strategy across European markets. It is this bottom-up lens that allows us to move from economy-wide narratives to concrete, submarket-level implications for space.

In the meantime, occupiers, investors and real estate professionals should treat AI exposure as an opportunity to prepare for, rather than minimise its implications. The evidence in this paper suggests that the cities, sectors, and firms that lean into AI-driven reorganisation – rather than delaying it – are the ones most likely to capture the next wave of productivity gains, business formation, and, in turn, real estate demand.

Contacts

Dennis Schoenmaker
Global Head of Forecasting &
Strategic Insight
dennis.Schoenmaker@cbre.com

Charles Zhang
Senior Analyst
Global Research
charles.zhang@cbre.com

© Copyright 2026. All rights reserved. This report has been prepared in good faith, based on CBRE's current anecdotal and evidence based views of the commercial real estate market. Although CBRE believes its views reflect market conditions on the date of this presentation, they are subject to significant uncertainties and contingencies, many of which are beyond CBRE's control. In addition, many of CBRE's views are opinion and/or projections based on CBRE's subjective analyses of current market circumstances. Other firms may have different opinions, projections and analyses, and actual market conditions in the future may cause CBRE's current views to later be incorrect. CBRE has no obligation to update its views herein if its opinions, projections, analyses or market circumstances later change.

Nothing in this report should be construed as an indicator of the future performance of CBRE's securities or of the performance of any other company's securities. You should not purchase or sell securities—of CBRE or any other company—based on the views herein. CBRE disclaims all liability for securities purchased or sold based on information herein, and by viewing this report, you waive all claims against CBRE as well as against CBRE's affiliates, officers, directors, employees, agents, advisers and representatives arising out of the accuracy, completeness, adequacy or your use of the information herein.