

Raising the Bar: The Impact of Generative AI on Experience Requirements in Job Postings

Hyunjae Kang, TaeYoung Kang, Sunham Kim

October 16, 2026

How does generative AI affect the demand for experienced workers? Using 10.4 million Korean job postings from July 2019 to January 2025, we estimate that ChatGPT's release increased required experience by 0.07 years per standard deviation of job-level AI exposure, and reduced the share of postings open to entry-level applicants. The effect remains statistically significant when IT occupations and IT-industry firms are excluded. Comparing each occupation with blue-collar jobs, we find the clearest increase among software developers, whose required experience rose by 0.62 years, and larger increases in occupations with high AI exposure.

Keywords: generative AI, job postings, experience requirements, entry-level jobs, Korea

JEL codes: J23, J24, J63, O33

1. Introduction

Since the advent of generative AI, employment and hiring of young workers have fallen in many jobs and industries exposed to the technology (Brynjolfsson, Chandar, and Chen 2026; Bank of Korea 2025). Generative AI can now do much of the work that new hires do in their first years. If employers respond by hiring more experienced workers, fewer postings should be open to entry-level workers. Whether employers now prefer more experienced workers for jobs more exposed to AI remains an open question.

In this paper, we study whether Korean employers raised required years of experience and opened fewer postings to entry-level applicants in jobs more exposed to generative AI, after ChatGPT’s release in November 2022. Among the postings a firm opens in a given month, we compare jobs that differ in exposure to generative AI. Then, we compare how the effects vary across occupations.

We construct a posting-level dataset from one of Korea’s largest job platforms, with 10.4 million postings dated July 2019 to January 2025. The dataset has three advantages. First, each posting states the experience it requires, which we convert to years. Because employers set these requirements before seeing any applicant, they measure whom employers are willing to hire, rather than an equilibrium outcome such as employment. Second, each posting lists job keywords, such as backend development, Java, Python, and firmware, which are much narrower than occupational groups. This granularity allows us to compare jobs that differ in exposure among the postings a firm opens in a given month, even within the same occupational group. Third, the platform covers a broad range of occupations, including blue-collar jobs, which enables us to compare exposed occupations with jobs that generative AI barely affects. We measure each job’s exposure to generative AI with the task-level scores of exposure to large language models (LLMs) from Eloundou et al. (2024), mapped to job keywords.¹ We score each posting by its first job keyword that has a score, which is usually the job itself.

In Korea, job postings for entry-level developers fell by about half between August 2022 and January 2025. After November 2022, when OpenAI released ChatGPT, the experience required in software-developer postings rose from an average of 3.2 years between July 2019 and July 2022 to 4.9 years in January 2025, an increase of 56%.

To estimate the impact of AI exposure across jobs, we employ a continuous difference-in-differences design that exploits variation in exposure among the jobs a firm posts in a given month. A firm-by-month fixed effect absorbs shocks to a firm in a given month, from its cost of credit to its reaction to the October 2022 credit-market shock. Our preferred specification adds occupational-group-by-month fixed effects, which control for monthly

¹Our crosswalk maps 1,920 job keywords, covering 99.8% of postings, to O*NET occupations, 800 of them through Korean occupation codes and the remaining 1,120 by direct matching. On the 681 keywords linked directly through those codes, direct matching gives scores that correlate at 0.85 with the code-linked scores.

shocks to each occupational group, such as IT and data. Thus, the 2021 developer boom and subsequent reversal are not attributed to exposure.

We find that the more exposed jobs began to require more experience after the release of ChatGPT. In our preferred specification, one standard deviation (SD) of exposure raises required experience by 0.07 years. This implies that, within the same firm and month, a job one SD more exposed came to require about 0.9 more months of experience, 3.4% of the pre-ChatGPT mean. The gap gradually increases after the initial release and becomes more pronounced in late 2023.

The increase in required experience comes with fewer postings open to entry-level workers. Per SD of exposure, the share of postings open to applicants without any experience falls by 0.77 percentage points. These changes suggest that employers may expect greater productivity gains from hiring experienced workers who are likely to direct generative AI more effectively.

Finally, within the same firms, we compare each occupation with blue-collar jobs, whose AI exposure is limited. The increase is clearest among software developers. Relative to blue-collar jobs, required experience in developer postings rose by 0.62 years, which is about 19% of the pre-release mean, and the share of developer postings open to entry-level workers fell by 7.2 percentage points. Marketing, IT infrastructure, and digital design show smaller increases. In logistics and customer service, the estimates are close to zero. Across occupations, the increase is larger where exposure is higher. Generative AI thus appears to raise the bar for entry most in the occupations where it can do more of the work that new hires used to do.

A growing number of papers estimate the impact of generative artificial intelligence on worker productivity (Noy and Zhang 2023; Brynjolfsson, Li, and Raymond 2025), employment and earnings (Humlum and Vestergaard 2025; Hosseini and Lichtinger 2026; Brynjolfsson, Chandar, and Chen 2026; Lodefalk et al. 2026; Lambert and Schindler 2026; Bank of Korea 2025; Korea Labor Institute 2026), and job postings (Audoly, Guerin, and Topa 2026; Iscenko and Curto Millet 2026; Westby, Modestino, and Cheng 2026). To the best of our knowledge, we are the first to estimate changes in required years of experience following ChatGPT’s release using a difference-in-differences design that compares jobs with different AI exposure posted by the same firm in the same month. Westby, Modestino, and Cheng (2026) estimate a decline in junior relative to senior software-developer postings using difference-in-differences. They also document rising mean experience requirements through a before-and-after comparison and a decomposition within and across job titles, rather than a difference-in-differences estimate of required years.

The rest of the paper is organized as follows. Section 2 describes the postings data and the exposure measure. Section 3 documents the hiring in the Korean job market around ChatGPT. Section 4 explains our empirical strategy, and Section 5 reports the

results. Section 6 concludes.

2. Data and Measurement

We use job postings data from one of the largest job search platforms in South Korea. The platform is popular among a wide range of employers, from large domestic and foreign companies such as Samsung, Hyundai Motor, and Applied Materials Korea to small and medium-sized enterprises and public institutions. The sample includes postings between July 2019 and January 2025. The data contain 10,413,271 postings, and 10,411,385 have firm and job identifiers. These postings come from 434,642 firms and have 17,543 distinct first job keywords. Each posting includes the posting date, experience and education requirements, employment type, pay, location, application dates, and ordered job keywords. Job keywords have two levels of description: first, an occupational group, such as IT development and data, and second, detailed keywords within that group, including programming languages and frameworks such as Java, Python, embedded, and firmware. We identify a job from a posting by its first job keyword. Our sample ends in January 2025 because the platform overhauled its experience labels on February 5, 2025.

Firms. We link each posting to its firm through the platform’s firm identifier. For 337,158 of the 434,688 firms behind all 10,413,271 postings (77.6%, covering 87% of postings), the platform’s company profiles give the firm’s industry, business description, founding date, and firm type. We further match 96.7% of these firms (75% of all firms) to business registration numbers, using the Korea Fair Trade Commission’s registry of mail-order businesses and the Financial Supervisory Service’s DART disclosure system, which allows the postings to be linked to administrative firm records. We use the firm identifier for the firm-by-month fixed effects and the industry to identify IT-industry firms.

Required experience. We treat entry-only postings and “no experience required” as zero years, a range from N to M years as $(N + M)/2$, N or more years as N , and N years or less or “entry or $N+$ years” as $N/2$. We treat “any experience” and experienced-worker postings without specific years as missing values. With this coding, required years are available for 85.5% of the analysis sample, with a mean of 2.3 years. In the analysis sample, 58.1% of postings explicitly welcome entry applicants, and 60.8% are open to applicants without experience once postings with plain “any experience” are included. After dropping postings without required years or exposure and singleton fixed-effect cells, the estimation sample has 6,998,390 postings, spanning 21 occupational groups.

AI exposure. We use the occupation-level LLM exposure scores of Eloundou et al. (2024), based on whether an LLM can halve the time to complete a task at the same quality.

Specifically, a task scores 1 if an LLM alone can do this, 0.5 if it requires additional software built on an LLM, and 0 otherwise. An occupation’s score is the average over all its tasks. Each posting receives the score of its first job keyword that has a score. We link 1,920 job keywords of the platform to O*NET occupations, so that each keyword takes its occupation’s score. When scoring a posting, we skip keywords that only repeat the occupational group’s name, such as “production” in the production group, because they do not identify a specific job. The scores are available for 99.8% of postings with firm and job identifiers. We standardize exposure across scored postings in the sample window before applying identifier and outcome restrictions.

3. The Korean Job Market Around the ChatGPT Shock


ChatGPT use has grown rapidly in South Korea.² Experience requirements rose and entry-level job postings declined in Korea after its release in November 2022. Figure 1 shows that average required experience remained stable at around 2.2 years. Monthly job posting volume also fluctuated around its mid-2021 level of about 170K, following a sharp decline and recovery during the initial COVID-19 outbreak in 2020. Following ChatGPT’s release, average required experience began to increase. Both total and entry-level job posting volumes remained relatively stable, with modest declines that were more pronounced for entry-level postings.

Coding postings required about 3.7 years of experience from July 2019 to July 2022, with little variation through the boom and subsequent reversal in coding posting volume. The requirement rose slightly to 4.1 years in December 2022, continued to rise, and reached 5.2 years in January 2025. Entry-level coding postings followed a similar trend to total postings until mid-2022 but then declined sharply.³ In contrast, the corresponding requirement for blue-collar jobs remained stable at around one year throughout, and entry-level blue-collar posting volume continued to follow the trend in all entry-level postings.


Experience requirements and posting volume could be correlated. Employers increase requirements when workers are abundant and relax them when they are scarce (Modestino, Shoag, and Ballance 2016, 2020), and these adjustments can be persistent (Hershbein and Kahn 2018). However, if a contraction in hiring pushed up experience thresholds in Ko-

²As of May 2025, South Korea had the largest number of paid ChatGPT subscribers outside the United States, and the number of weekly ChatGPT users in the country had grown by a factor of 4.5 over the preceding year (Tan 2025).

³We count a posting toward coding in proportion to the share of its job keywords that are stable coding keywords, such as “Java”, “Python”, “backend/server development”, and “frontend”. Postings for software developers fall into this category. This weighting differs from the first-keyword definition in the Introduction and Table 2, under which required experience in developer postings rose from 3.2 to 4.9 years over the same period and stayed around 0.7 years in blue-collar postings. Stable keywords are those whose usage changed little when the platform expanded its keyword menu in September 2021, which keeps the series comparable over time.



A. Average required experience by Occupational Group



B. Total and Entry-level Job Postings

FIGURE 1. Experience requirements and job postings, July 2019 to January 2025

Notes. Panel A shows the monthly mean required experience by occupational group. The means are weighted by stable job-keyword shares. Values in parentheses are for January 2025. Panel B shows monthly job posting counts, with each series normalized to 100 in October 2022.

rea, the sharp decline in hiring during the initial pandemic shock in 2020 would have been accompanied by a significant and persistent increase in required experience. Moreover, the coding hiring boom in 2021 did not come with a decline in required experience.

Our descriptive plots, however, also suggest potential macroeconomic confounding. Postings declined slightly, and required experience increased slightly, just before ChatGPT’s release. Sales job posting volume followed the trend in total postings, but required experience increased. These changes could reflect macroeconomic conditions, such as monetary tightening. The Bank of Korea held its base rate at 0.50% until July 2021 and then raised it ten times between August 2021 and January 2023, reaching 3.50%. The rate remained at that level until October 2024 (Figure A1). These rate increases could have had lagged effects, which may be reflected in the trends. Moreover, the U.S. Federal Reserve’s repeated increases in the federal funds rate between March 2022 and July 2023 may also have affected Korean capital markets and, in turn, firms’ hiring decisions. To address these issues, we estimate the impact of the ChatGPT shock using variation in AI exposure within the same firm and month in the empirical analysis.

4. Empirical Strategy

Our empirical strategy consists of three specifications. First, to estimate the overall effect of AI exposure across occupations, we estimate its effect on required experience in a continuous difference-in-differences framework. Second, we trace the timing of the effect with an event study. Third, to estimate heterogeneous effects by occupation, we estimate a difference-in-differences equation with a binary treatment separately for each occupation, with blue-collar jobs as the comparison group.

Define y_{fijt} as required experience in years for firm f ’s posting i for job type j in month t ; z_i as standardized AI exposure of posting i in occupational group g . We estimate a continuous difference-in-differences regression equation.

$$(1) \quad y_{fijt} = \rho z_i + \beta(z_i \times \text{Post}_t) + \alpha_j + \delta_{ft} + \lambda_{g(i),t} + \varepsilon_{fijt}.$$

where Post_t equals one from December 2022 onward. The fixed effects are by job (α_j), firm-by-month (δ_{ft}), and occupational-group-by-month ($\lambda_{g(i),t}$), where $g(i)$ denotes the posting’s occupational group.

Postings for the same job can differ in exposure because the job is defined by the first keyword, while exposure comes from the first keyword with a score. When the first keyword has no score or repeats the occupational-group name, exposure is taken from a later keyword, and that keyword varies across postings. The term ρz_i absorbs any pre-ChatGPT difference in required experience between them, so that β measures only the change after launch. Standard errors are two-way clustered by firm and job.

Under parallel counterfactual changes in required experience across exposure levels, conditional on these fixed effects, β identifies the differential effect of ChatGPT's release per standard deviation of exposure. We also report three alternative specifications: (i) one with only job and month fixed effects, (ii) one with firm-by-month but without occupational-group-by-month fixed effects, and (iii) one that adds firm-by-job fixed effects, which absorb the job fixed effects. We do not use firm-by-job fixed effects in the preferred specification because they identify the effect only from firm-job pairs that post both before and after ChatGPT's release. If generative AI leads firms to stop posting entry-level openings for a job, those pairs drop out of the comparison, so the estimate would exclude variation of interest.

For the event study, we replace the single post-ChatGPT interaction with period-specific exposure interactions while keeping the main exposure effect:

$$(2) \quad y_{fijt} = \rho z_i + \sum_{k \neq \text{ref}} \beta_k (z_i \times \mathbf{1}[t \in k]) + \alpha_j + \delta_{ft} + \lambda_{g(i),t} + \varepsilon_{fijt}.$$

We pool from July 2019 to May 2021 as a reference period. Each subsequent unit is three months, starting June 2021. The first post-launch unit is December 2022–February 2023. The final unit, December 2024–January 2025, contains two months. We use the same design for the entry-eligibility event studies.

Finally, we estimate heterogeneous effects by occupation. For each occupation o , we estimate

$$(3) \quad y_{fijt} = \theta_o (D_j^o \times \text{Post}_t) + \alpha_j + \mu_f + \tau_t + \varepsilon_{fijt}$$

on postings in occupation o and in blue-collar jobs, where D_j^o equals one if job j belongs to occupation o and zero if it is a blue-collar job. The job fixed effect α_j absorbs permanent differences between jobs, including the level of D_j^o . The firm fixed effect μ_f absorbs permanent differences between firms, and the month fixed effect τ_t absorbs shocks common to all postings. The coefficient θ_o measures how much more required experience rose in occupation o than in blue-collar jobs after ChatGPT's release. We estimate the same equation for entry eligibility, and standard errors are two-way clustered by firm and job.

5. Results

Required experience rises more in jobs more exposed to AI. Table 1 reports the estimates across fixed effects and samples. With job and month fixed effects, the estimate is 0.113 years per exposure standard deviation. With firm-by-month effects it is 0.133 years. Adding occupational-group-by-month effects yields our preferred estimate of 0.073 years ($p =$




FIGURE 2. Event-study coefficients, preferred specification

Notes. All eligible postings; estimation sample $N = 6,998,390$. Coefficients are required years per exposure standard deviation relative to July 2019–May 2021. Bars show 95% confidence intervals with standard errors clustered by firm and job. Hollow markers are not significant at 5%. The dashed line separates the final pre-launch bin from December 2022–February 2023. The model includes the exposure main effect and job, firm-by-month, and occupational-group-by-month fixed effects. The final bin covers December 2024–January 2025. A joint test of the six pre-launch coefficients gives $p = 0.371$.

0.002), based on 6,998,390 postings.⁴ This corresponds to 0.87 months or 3.4% of the pre-launch mean among regression-eligible postings. For comparison, Westby, Modestino, and Cheng (2026) find that mean required experience in US software-developer postings rose by 0.34 years, or 6%, in the year after ChatGPT’s release. The two magnitudes are not directly comparable. There is a change over time in one occupation without a comparison group, but ours compares jobs within the same firm and month across all occupations, net of shocks to each occupational group. Measured in the same way, mean required experience in our software-developer postings rose from 3.24 years in 2022Q3 to 3.92 years in 2023Q3, an increase of 21%. This increase is considerably larger than in the United States, both in years and in percentage terms, from a lower starting level. Relative to blue-collar jobs posted by the same firms, required experience in developer postings rose by 0.62 years after launch (Table 2, column 1), so the rise is not shared by blue-collar jobs.

The effect rises gradually after the launch. Figure 2 reports our preferred event study specification. Post-launch coefficients are positive and first become significant at 5% in September–November 2023, remaining significant through the end of the sample. The point estimates also increase before launch in the second half of 2022.

⁴A placebo test with a fake post date of July 2021 gives -0.003 years (Table A1, column 1).

TABLE 1. AI exposure, experience requirements, and entry-level eligibility

	(1) Job + month	(2) Firm- month	(3) Preferred	(4) Firm- job	(5) Non-IT	(6) Short window
<i>A. Required experience (years): specification-specific samples</i>						
Exposure $\times$ Post	0.1128***	0.1333***	0.0726***	0.0516*	0.0551**	0.0531*
(standard error)	(0.0220)	(0.0199)	(0.0235)	(0.0290)	(0.0236)	(0.0274)
Observations	8,880,085	6,998,390	6,998,390	6,464,223	6,031,755	2,903,646
<i>B. Required experience (years): identical firm-by-job sample</i>						
Exposure $\times$ Post	0.1216***	0.1400***	0.0768***	0.0516*	–	–
(standard error)	(0.0263)	(0.0215)	(0.0256)	(0.0290)	–	–
Observations	6,464,223	6,464,223	6,464,223	6,464,223	–	–
<i>C. Entry-level outcomes (percentage points): all classified labels</i>						
Broader eligibility	–	–	-0.7742***	-0.4152*	-0.7639***	-0.5718***
(standard error)	–	–	(0.2562)	(0.2490)	(0.2777)	(0.2041)
Explicit entry-level	–	–	-0.5016*	-0.2717	-0.4445	-0.3677*
(standard error)	–	–	(0.2695)	(0.2284)	(0.3006)	(0.2143)
Observations	–	–	8,315,761	7,704,096	7,192,984	3,460,952
Job FE	Yes	Yes	Yes	No	Yes	Yes
Month FE	Yes	No	No	No	No	No
Firm $\times$ month FE	No	Yes	Yes	Yes	Yes	Yes
Occupational group $\times$ month FE	No	No	Yes	Yes	Yes	Yes
Firm $\times$ job FE	No	No	No	Yes	No	No

Notes. All coefficients are on Exposure $\times$ Post. Post starts December 2022. Column (6) uses December 2021–November 2023; other columns use July 2019–January 2025. Exposure normalization is held fixed. Panel B holds exactly 6,464,223 row identities fixed in columns (1)–(4). Non-IT necessarily changes the population. Full analysis-sample shares are 58.1% for explicit entry-level mentions and 60.8% for broader eligibility. Panel C uses linear-probability models, including classified labels without numeric years. Broader eligibility includes explicit entry-level mentions and plain unrestricted experience; explicit entry-level counts only explicit mentions. Experienced-only labels are ineligible. Both outcomes share the sample within each column. Unestimated cells are shown as dashes. Standard errors are clustered by firm and job; p -values use a normal reference. Counts exclude singleton fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.




FIGURE 3. Entry-level eligibility event studies, preferred fixed effects

Notes. Coefficients are percentage points per exposure standard deviation relative to July 2019–May 2021, as in Figure 2. Panel A includes postings explicitly welcoming entry-level applicants and plain unrestricted-experience postings; Panel B counts explicit entry-level mentions. Both models retain 8,315,761 postings with classified labels, retain the exposure main effect, and include job, firm-by-month, and occupational-group-by-month fixed effects. Standard errors are clustered by firm and job. Bars show 95% confidence intervals; hollow markers are not significant at 5%. The dashed line separates the final pre-launch bin from December 2022–February 2023. The final bin covers December 2024–January 2025. Joint tests of the six pre-launch coefficients give $p = 0.294$ in Panel A and $p = 0.616$ in Panel B.

Exposed jobs also opened fewer postings to entry-level applicants. In the preferred specification, 1 SD of exposure lowers the share of postings open to applicants without experience by 0.77 percentage points after launch (Table 1, Panel C, column 3), against a sample share of 60.8%. The share of postings that explicitly express interest in entry-level applicants falls by 0.50 percentage points (significant at 10%). Figure 3 reports the event study, and the six pre-launch coefficients are jointly insignificant ($p = 0.29$ and $p = 0.62$). As with required experience, broader eligibility starts to fall in the second half of 2022, whereas explicit entry-level mentions decline only after launch. This decline is consistent with evidence that employers have shifted away from junior hiring in exposed jobs. In Sweden, the employment of workers aged 22 to 25 years in exposed occupations fell by 5.5% relative to less exposed occupations (Lodefalk et al. 2026). Required experience also rose among postings seeking experienced candidates, by 0.079 years per standard deviation of exposure (Table A1), so the rise in mean required experience is not only the result of fewer postings open to entrants.

Software developers show the clearest increase in required experience. Relative to blue-collar jobs, required experience in developer postings rose by 0.62 years after launch (Table 2), which is approximately 7 months or 19% of the pre-launch mean of 3.18 years. The share of developer postings open to applicants without experience fell by 7.2 percentage points from 43.2%. Both placebo estimates are close to zero. These results are in line with evidence from the United States, where postings for software developers requiring 3 or fewer years of experience fell 14 to 15% relative to those requiring 4 or more after November 2022 (Westby, Modestino, and Cheng 2026). In Korea, the share of software-development postings open to newcomers fell from 53.5% in 2022 to 37.4% in 2024 (Korea Labor Institute 2025). Our estimates show that this shift holds relative to blue-collar jobs posted by the same firms.

Other occupations exposed to AI show smaller increases. Required experience rose by 0.44 years in marketing, 0.43 years in IT infrastructure, and 0.27 years in digital design, with placebo estimates close to zero, and the share of postings open to applicants without experience fell by 3.0 to 5.6 percentage points in these occupations. Physical design and video show smaller increases, of 0.16 and 0.21 years. Data analysis has the largest point estimate (0.76 years), but its placebo estimate points to movements before launch, and the estimates for bookkeeping, clerical and office jobs, and construction sites also fail the placebo test. Required experience did not rise meaningfully in sales, logistics, HR, customer service, translation and content, education, or among tax accountants. Across the 17 occupations, the estimates correlate at 0.53 with mean exposure for required experience and at -0.35 for entry eligibility. This gradient mirrors the continuous estimates and the concentration of employment declines for young workers in the most exposed occupations in the United States (Brynjolfsson, Chandar, and Chen 2026). It is also hard to

TABLE 2. Required experience and entry eligibility after ChatGPT's release, by occupation relative to blue-collar jobs

Occupation	Jobs	Mean exposure	Required experience (years)			Entry eligibility (pp)		
			Pre-launch mean	(1) Effect	(2) Placebo	Pre-launch mean	(3) Effect	(4) Placebo
Software developers	21	0.89	3.18	0.62*** (0.07)	-0.02 (0.06)	43.2	-7.2*** (1.1)	0.2 (0.7)
Bookkeeping	9	0.78	2.04	0.19*** (0.07)	0.10*** (0.03)	57.1	-2.2*** (0.4)	-1.1 (0.8)
Data analysis	3	0.75	4.14	0.76*** (0.16)	-0.15* (0.08)	33.6	-6.7*** (0.9)	-3.0*** (1.1)
Translation, content	7	0.75	2.16	0.08 (0.09)	-0.03 (0.06)	55.1	-0.4 (1.1)	-0.1 (1.0)
Digital design	12	0.66	2.09	0.27*** (0.06)	0.05 (0.06)	51.5	-4.7*** (0.4)	-0.5 (0.6)
Tax accountants, CPAs	2	0.60	4.29	0.03 (0.05)	0.06 (0.04)	34.7	-1.5*** (0.5)	0.2 (0.6)
Marketing	7	0.58	3.77	0.44*** (0.04)	-0.03 (0.07)	37.0	-3.0*** (0.6)	-0.7 (0.9)
Customer service	6	0.58	0.75	0.02 (0.07)	-0.06 (0.05)	80.9	-0.1 (0.6)	1.9* (1.1)
Clerical and office	5	0.57	0.45	0.06*** (0.02)	0.10*** (0.02)	86.2	-1.6*** (0.5)	-2.4*** (0.5)
HR and general affairs	10	0.55	1.87	0.02 (0.08)	-0.10 (0.10)	63.3	0.5 (0.9)	0.5 (1.1)
Education	5	0.53	0.78	0.04*** (0.02)	0.04 (0.03)	71.9	-0.8 (0.7)	-0.3 (0.6)
IT infrastructure	5	0.51	2.56	0.43*** (0.13)	0.03 (0.08)	60.3	-5.6*** (0.7)	-1.0 (1.0)
Logistics, purchasing	6	0.50	1.41	-0.02 (0.03)	-0.01 (0.06)	74.9	-0.1 (0.5)	-1.2 (1.2)
Sales	8	0.49	2.24	0.09 (0.05)	-0.11*** (0.04)	64.3	-0.4 (0.5)	0.1 (0.5)
Video	6	0.49	1.23	0.21*** (0.05)	0.10* (0.06)	65.4	-3.9*** (0.9)	-0.5 (1.0)
Physical design	12	0.47	2.99	0.16*** (0.07)	-0.07 (0.07)	47.9	-4.6*** (0.7)	0.5 (1.2)
Construction sites	5	0.40	2.77	0.10*** (0.04)	-0.19*** (0.05)	49.5	-2.8** (1.2)	2.2 (1.4)
Blue-collar jobs	246	0.11	0.68			80.9		
Corr. with exposure				0.53			-0.35	

Notes. Columns (1) and (3) report the coefficient on the occupation indicator interacted with Post (December 2022 onward), estimated on postings in the occupation and in blue-collar jobs with job, firm, and month fixed effects. Columns (2) and (4) use postings through November 2022 with a placebo Post date of July 2021. A posting's occupation is that of its first job keyword; blue-collar jobs are keywords in no occupation with exposure of at most 0.2. Entry eligibility means open to applicants without experience. Pre-launch means cover July 2019 to November 2022. The last row correlates mean exposure with columns (1) and (3) across the 17 occupations. Standard errors, clustered by firm and job, may understate uncertainty in occupations with few jobs. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.

reconcile with a general cyclical rise in requirements, which would have raised requirements in less exposed white-collar occupations as well. Raw requirements in sales did rise (Figure 1A), but relative to blue-collar postings of the same firms, they did not rise in sales or logistics (Table 2).

5.1. Robustness Checks

Hiring cycle and monetary tightening. Firm-by-month fixed effects absorb shocks common to a firm’s postings, including its cost of credit. The exposure gradient did not move during the Bank of Korea’s first five hikes (August 2021–May 2022; Figure 2), and it kept widening while the base rate was held at 3.50% from January 2023 to September 2024, by 0.031 years per SD per year. Controlling for each job’s pre-launch boom, cooling and rate sensitivity leaves the estimate at 0.082, compared with 0.080 on the same sample (Table A1). Dropping the IT development/data occupational group and IT-industry firms, which removes the 2021 developer boom and its reversal, leaves an estimate of 0.055 years, based on 6,031,755 postings.

Firm-by-job fixed effects. Replacing job fixed effects with firm-by-job fixed effects, while retaining firm-by-month and occupational-group-by-month effects, gives an estimate of 0.052 years, based on 6,464,223 postings. The estimate remains positive and is significant at 10%, but not at 5%. On identical observations, adding firm-by-job fixed effects reduces the estimate from 0.077 to 0.052 years (Table 1, Panel B).

Pre-launch event-study coefficients. The six pre-launch coefficients in Figure 2 are not jointly distinguishable from zero relative to the pooled reference period. The late-2022 point estimates rise before launch. Non-rejection of the joint test, however, doesn’t establish parallel trends.

6. Conclusion

We study whether employers began to demand more experience for jobs that generative AI can do. Using South Korean job posting data, we compare jobs with different exposure to generative AI posted by the same firm in the same month, using blue-collar jobs as a comparison group. We find a significant increase of required experience for jobs with higher AI exposures after ChatGPT’s release. Exposed jobs also opened fewer postings to entry-level applicants and asked for more experience even in postings that already required experience.

Our estimates indicate changes in what the demand side asks for, but we do not model the firm’s decision behind it. We leave it for future research.

References

- Audoly, Richard, Miles Guerin, and Giorgio Topa. 2026. “Do Job Postings Show Early Labor-Market Effects of AI?” Liberty Street Economics, Federal Reserve Bank of New York, May 14. <https://doi.org/10.59576/lse.20260514>.
- Bank of Korea. 2025. “AI Diffusion and the Contraction of Youth Employment: Seniority-Biased Technological Change.” BOK Issue Note 2025-30, Bank of Korea. By Jinsu Han and Samil Oh. October 30. In Korean.
- Brynjolfsson, Erik, Bharat Chandar, and Ruyu Chen. 2026. “Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.” Working paper, Stanford Digital Economy Lab, August 2026 version.
- Brynjolfsson, Erik, Danielle Li, and Lindsey Raymond. 2025. “Generative AI at Work.” *Quarterly Journal of Economics* 140 (2): 889–942.
- Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2024. “GPTs are GPTs: Labor Market Impact Potential of LLMs.” *Science* 384 (6702): 1306–1308.
- Hershbein, Brad and Lisa B. Kahn. 2018. “Do Recessions Accelerate Routine-Biased Technological Change? Evidence from Vacancy Postings.” *American Economic Review* 108 (7): 1737–1772.
- Hosseini, Seyed M. and Guy Lichtinger. 2026. “Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data.” Working paper, May 2026 version.
- Humlum, Anders and Emilie Vestergaard. 2025. “Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI.” NBER Working Paper 33777, National Bureau of Economic Research. May 2025, revised March 2026. Previously circulated as “Large Language Models, Small Labor Market Effects”.
- Iscenko, Zanna and Fabien Curto Millet. 2026. “Looking for the Ladder: Is AI Impacting Entry-Level Jobs?” Economic Innovation Group, January.
- Korea Labor Institute. 2025. “Recent Changes in the Number and Work of Software Developers.” *Monthly Labor Review* (249): 75–89. By Sanghoon Ji. December 2025. In Korean.
- Korea Labor Institute. 2026. “The Effects of the Diffusion of AI Technology on Employment.” *Monthly Labor Review* (251): 47–65. By Jiyeon Jang. February 2026. In Korean.
- Lambert, Peter John and Yannick Schindler. 2026. “The Broken Ladder: AI, Remote Work, and Early-Career Hiring.” CAGE Working Paper 808, University of Warwick. May 2026. Also SSRN Working Paper 6787638.
- Lodefalk, Magnus, Lydia Löthman, Michael Koch, and Erik Engberg. 2026. “Same Storm, Different Boats: Generative AI and the Age Gradient in Hiring.” Working Paper (Economics) 2/2026, Örebro University School of Business.
- Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2016. “Downskilling: changes in employer skill requirements over the business cycle.” *Labour Economics* 41: 333–347. <https://doi.org/10.1016/j.labeco.2016.05.010>. SOLE/EALE conference issue 2015.
- Modestino, Alicia Sasser, Daniel Shoag, and Joshua Ballance. 2020. “Upskilling: Do Employers Demand Greater Skill When Workers Are Plentiful?” *Review of Economics and Statistics* 102 (4): 793–805.
- Noy, Shakked and Whitney Zhang. 2023. “Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence.” *Science* 381 (6654): 187–192.
- Tan, Kwan Wei Kevin. 2025. “OpenAI Is Opening a New Office in South Korea to Serve Its Boom-

ing Consumer Base in the Country.” *Business Insider*. <https://www.businessinsider.com/openai-new-office-south-korea-booming-consumer-base-2025-5>. May 26, 2025.

Westby, Samuel, Alicia Modestino, and Peiran Cheng. 2026. “Generative AI and the Redefinition of Entry-Level Software Work.” IZA Discussion Paper 18723, IZA. June 2026.

Appendix A. Additional Results

TABLE A1. Robustness of the main estimate and margins of the increase

<i>A. Robustness of the main estimate: required experience (years)</i>							
	(1)	(2)	(3)	(4)	(5)	(6)	(7)
	Main	Cycle sample	Cycle controls	Skill sample	Skill controls	First tag	Most exposed
Exposure × Post	0.073*** (0.023)	0.080*** (0.026)	0.082*** (0.026)	0.073*** (0.023)	0.083*** (0.023)	0.070*** (0.023)	0.039** (0.017)
Placebo: Exposure × Post	−0.003 (0.035)				0.006 (0.031)	−0.004 (0.035)	−0.005 (0.020)
Boom size × Post			0.108** (0.047)				
Cooling size × Post			−0.086** (0.040)				
Rate sensitivity × Post			0.101*** (0.037)				
Pre-launch experience × Post					−0.085** (0.036)		
Degree share × Post					0.003 (0.032)		
Posted salary × Post					0.048** (0.021)		
Observations	6,998,390	5,530,544	5,530,544	6,934,379	6,934,379	6,933,576	6,998,390
<i>B. Margins of the increase</i>							
	(1) Years, experienced postings	(2) Requires 3+ years (pp)	(3) Requires 5+ years (pp)				
Exposure × Post	0.079*** (0.029)	0.63*** (0.21)	0.66*** (0.22)				
Placebo: Exposure × Post		0.18 (0.29)	0.04 (0.36)				
Observations	3,398,894	6,998,390	6,998,390				

Notes. All regressions follow our preferred specification (Table 1, column 3), with the exposure main effect and job, firm-by-month, and occupational-group-by-month fixed effects; Post starts December 2022. Placebo rows use postings through November 2022 with a fake post date of July 2021. In Panel A, columns (2) and (3) use the 596 jobs with at least 500 postings and add each job's boom (posting growth from 2020H2 to 2021H2), cooling (from March–August 2021 to October 2021–November 2022), and pre-launch rate sensitivity of required experience. Columns (4) and (5) add each job's pre-launch required experience, four-year degree share, and log posted salary. All controls are standardized and interacted with Post. The cycle measures are built from the placebo window, so no placebo is reported for them. Columns (6) and (7) measure exposure by the first job keyword and by the most exposed keyword. In Panel B, column (1) uses postings that require experience, and columns (2) and (3) are linear probability models in percentage points. Standard errors, two-way clustered by firm and job, are in parentheses. Observations exclude singleton fixed effects. *** $p < 0.01$, ** $p < 0.05$, * $p < 0.10$.




FIGURE A1. Monetary Policy Changes in South Korea and U.S.