

Generative AI and the Reconfiguration of Skill Demand in Chinese Job Postings

Yuan Dong

dyuan631@snnu.edu.cn

Shaanxi Normal University

Research Article

Keywords: generative AI, online vacancies, labour demand, task reallocation, skill combinations, entry requirements

Posted Date: October 21st, 2026

DOI: <https://doi.org/10.21203/rs.3.rs-11173388/v1>

License:   This work is licensed under a Creative Commons Attribution 4.0 International License.

[Read Full License](#)

Additional Declarations: No competing interests reported.

Generative AI and the Reconfiguration of Skill Demand in Chinese Job Postings

Yuan Dong

International Business School, Shaanxi Normal University, Shaanxi China

Email: dyuan631@snnu.edu.cn

Address for correspondence:

Yuan Dong,

International Business School, Shaanxi Normal University, Shaanxi, China

Email: dyuan631@snnu.edu.cn

Phone: 86-17321731631

Abstract

Generative artificial intelligence (AI) can change the tasks employers assign to jobs and the qualifications they require of applicants. Using Chinese online job postings from 2018 to 2025, we link vacancy descriptions to a fixed ontology of 2,083 tasks and construct four complementary measures of employer demand: routine tasks, task reallocation, within-job task combinations, and entry requirements. We estimate models at the province-month, job-title, and task levels and compare the corresponding samples and inference procedures. Routine-task demand shows no robust average change in the main specification, whereas exposure-related changes in task allocation are visible across task groups. Within-job task-combination distance increases, although the composite combination index is less robust to wild-cluster inference. Educational requirements do not rise consistently; requested work experience increases in the task-level analysis. Placebo and pre-trend diagnostics limit a uniform causal interpretation, particularly for some reallocation and combination results. Job advertisements measure employer requirements rather than realised educational expenditure or on-the-job training. The findings therefore describe demand-side incentives for acquiring adaptable and complementary skills, rather than observed human capital investment.

Keywords: generative AI; online vacancies; labour demand; task reallocation; skill combinations; entry requirements

JEL classification: J23; J24; O33

1 Introduction

Large language models extend the range of computationally exposed work to writing, analysis, programming and other cognitive tasks. Exposure measures technological feasibility, rather than the disappearance of an occupation^[1]. Experiments document productivity gains in professional writing^[2] and software development^[3], while workplace evidence suggests particularly large assistance effects for less experienced customer-service staff^[4]. Knowledge-work experiments also find that performance depends on the fit between the assignment and the technology^[5]. Taken together, these findings motivate a direct examination of what employers ask new hires to do and know as generative artificial intelligence (AI) spreads.

Human capital theory regards schooling and training as present investments that generate later returns^[6]. Earnings depend on both formal education and accumulated experience^[7], while the remaining working life changes the incentive to invest at different ages^[8]. The central question in this paper is therefore whether a change in the relative value of job tasks alters employer signals concerning qualifications acquired before labour-market entry and capabilities updated during employment. Rapid adoption of general-purpose AI tools^[9] makes this question timely, but adoption surveys cannot substitute for observed vacancies.

This distinction between the timing of formation and the timing of use motivates the paper’s investment terminology. Formal schooling can provide broad cognitive and disciplinary foundations well before a future job is known. Experience gained in employment may be specific to a task, a set of clients or an organisation, even while some abilities remain transferable. When the mix of tasks changes, the expected value of each route depends on whether capabilities can be recombined and whether the entry screen prevents candidates from demonstrating those abilities. Our empirical strategy does not assign a worker to either a “schooling

investor” or a “training investor”. It asks whether employers change the mix of requirements that could shape future incentives for those choices. The same person can undertake both routes at different points in a career, and a static advertisement cannot reveal the optimal timing or intensity of either investment. This interpretation stays close to the life-cycle mechanisms in^[6–8] while respecting the observational limits of vacancies.

We distinguish four related margins of employer demand. Routine-task demand records whether established task requirements persist. Task reallocation measures the redistribution of demand across a fixed set of tasks. Within-job combinations describe how tasks are bundled into vacancies. Entry requirements record education and prior experience. The first three measure aspects of job content, whereas the fourth measures screening. Lazear’s skill-weights account suggests that even portable skills may have job-specific value when combined in different proportions^[10]. None of these vacancy measures records a worker’s educational expenditure, training participation or realised investment return.

The study maps approximately 5.698 million standardised Chinese online vacancy records from 2018–2025 to a fixed ontology of 2,083 tasks. It constructs demand outcomes at the province-month, province-month-job-title, and province-month-task levels, retaining the denominator and valid sample appropriate to each measure. The matching design makes it possible to connect an employer’s task allocation to historically comparable task bundles and advertised qualifications. Earlier work documents changes in AI-specific hiring and other skill requirements^[11]; we examine additional dimensions of the recruitment structure without treating them as observations of actual worker investments.

The distinction between demand and investment also clarifies why a vacancy is an appropriate, but limited, observation for this question. A vacancy describes a position that an employer plans to fill; its responsibilities indicate the tasks the employer expects somebody to perform, and its requirements indicate the attributes used to screen candidates. These are observable before a match is formed and can therefore react relatively promptly to revised expectations about technology. A realised wage, in contrast, reflects the eventual match and other bargaining and selection processes. We do not assert that a listed vacancy was filled, that its author used an AI system, or that a listed requirement describes every incumbent in the firm. Our

interpretation is conditional on that recruitment stage. The main contribution is to disaggregate a potentially broad statement about “human capital investment” into four employer-side signals with explicit, different units of observation. This narrower empirical object allows a sharper comparison of task content and screening without requiring us to reconstruct workers’ unobserved choices from employer text.

The evidence is heterogeneous across margins. The main specifications do not identify a stable average shift in routine demand. Task shares differ across exposure groups following diffusion, and within-job task-combination distance increases; the latter result survives wild-cluster inference in the valid combination sample. A task-level design finds a positive association with experience requirements but no comparable stable increase in average years of schooling. Some lead and pre-trend diagnostics caution against attributing every change to a single December 2022 discontinuity. The paper consequently interprets the results as differentiated demand-side incentives for continuing adaptation, subject to explicit limits on causal and individual-level interpretation.

A second contribution concerns measurement consistency. When a current posting is compared with earlier postings bearing the same title, the meaning of a “different” bundle depends on a stable task coordinate system and a well-defined historical reference set. The fixed ontology makes a task in one year comparable to its counterpart in another, while the rolling title-by-province profile limits comparisons between substantially different occupations. The reallocation series, by contrast, deliberately uses the entire task support and a prior-year same-month comparison. Routine-task shares and entry requirements have still other denominators. These design choices are not interchangeable technical details: they determine which change a coefficient can represent. We retain the distinct scales and samples throughout rather than collapsing several heterogeneous outcomes into a single index of human capital investment.

2 Literature, mechanism and empirical implications

2.1 Human capital formation over the working life

Becker’s investment framework^[6], Mincer’s education–experience specification^[7], and Ben-Porath’s life-cycle model^[8] jointly imply that skill formation need not end when a worker enters the labour market. The long-run relative supply of educated workers and technical demand shapes the payoff to formal

qualifications^[12]. On the job, task assignment and learning can generate task-specific human capital^[13], while different weights on otherwise general skills can form a valuable firm-specific bundle^[10]. A technology that changes task weights may therefore change both the value of pre-entry credentials and the incentive to update skills later.

Vacancy text is a demand-side lens on this process. Deming and Kahn^[14] document variation in professional skill requirements across firms and markets; Acemoglu and co-authors^[11] use online vacancies to study AI hiring. We interpret education and experience requirements as entry screening and task content as a signal of work to be performed. An advertised threshold does not reveal whether a candidate has invested more, whether an incumbent receives training, or whether employers eventually pay a higher wage for the stated skill. This distinction governs every hypothesis and conclusion below.

There is also an important distinction between the level and composition of employer requirements. Suppose the number of postings for a province rises while the relative prevalence of an exposed task declines. The decline in its share would be compatible with an increase, decrease or no change in the number of postings containing that task. Similarly, a larger within-posting distance does not mean that a posting includes more tasks or that every task has become more difficult. It means that, relative to an appropriately defined recent history, the vector of tasks in current vacancies differs. Schooling years and experience years are conditional means among vacancies with interpretable numeric requirements; an increase in the share of unrestricted vacancies can change the population to which those means apply. Reporting the four outcomes together is informative precisely because their meanings do not collapse into one scalar response.

2.2 Automation, augmentation and task reallocation

Occupation-level assessments of computerisation^[15] identify heterogeneous exposure, but task-based accounts distinguish substitution of particular activities from the complementary tasks that remain^[16,17]. Automation can displace some work while new tasks restore other forms of labour demand^[18,19]. Assessments of generative AI likewise separate potential task exposure from actual employment change^[1,20]. Thus the average response of routine-task demand has no theoretically predetermined sign: productivity may reduce

recruitment for an established task, expand output sufficiently to maintain it, or shift the activity toward monitoring and verification.

Writing experiments^[2], customer-service evidence^[4] and work on AI as an aid to prediction^[21] are consistent with effects that vary by task and worker experience. A rise in efficiency alone does not establish an increase in hiring. Accordingly, our first implication is to estimate routine-task share and intensity separately from the distribution of task-share gains and losses. An imprecise average estimate should be interpreted with its confidence interval, rather than as proof of a strict zero effect. Main Tables 4 and 5 correspond to these two margins.

Chinese vacancy studies report changes in AI-related skills and wages^[22] and differentiated recruitment around the release of ChatGPT^[23]. Studies of employees also indicate that workplace diffusion and labour-market adjustment need not occur simultaneously^[24]. Results based on listed-firm vacancies suggest increasing education and experience thresholds in a different sample^[25]. Such findings make sampling, outcome definition and observation level essential when comparing studies; disagreement over an average schooling outcome in separate samples cannot be resolved by comparing significance stars.

Changing task content can raise demand for nonroutine and interactive abilities^[26,27], while routine-biased change is relevant for polarisation^[28]. Historical advertisement data show that the content of jobs can evolve within an occupation^[29]. These observations motivate the measurement of combinations within a stable task ontology, in addition to changes in the distribution of tasks across vacancies.

The literature further suggests that technological capability and realised organisational adjustment can be separated in time. An exposure index represents tasks for which a model might be useful. Employers still must decide how to verify output, protect data, allocate responsibility and redesign coordination before that potential becomes a change in a job description. New positions can then contain a mixture of retained, automated and augmented activities. Consequently, an immediate post-release coefficient is only one summary of a process that may have begun with earlier tools and expectations and may continue after widespread availability. This timing distinction motivates the stage-specific results and lead specifications rather than serving as an ex post

explanation for inconvenient tests. Where a lead is informative, we explicitly reduce the strength of a point-in-time causal interpretation.

This is also why employment totals alone cannot answer the present question. A firm might recruit the same number of people while altering their responsibilities, or redistribute tasks among positions while maintaining its stated number of openings. In either case, a demand measure based solely on posting counts would miss a change in the capabilities that new workers are expected to supply. Conversely, an increase in openings can coincide with more standardised tasks and unchanged entry thresholds. The four margins therefore describe different slices of the same recruiting process and furnish an internal comparison: whether established work persists, whether the task portfolio shifts, whether a vacancy bundles tasks differently and whether selection criteria move. We use their joint pattern to discuss a possible transition toward continuous capability renewal, while avoiding the stronger claim that any one marginal coefficient proves a completed change in investment behaviour.

2.3 Skill bundles, ongoing adaptation and entry screening

The value of a skill bundle is not necessarily the sum of isolated skill values. A firm may benefit from a distinctive weighting of transferable competences^[10], and vacancy-based evidence finds different bundles across employers^[14]. Our second implication concerns the distance between the task vector of a new posting and the recent task profile of postings with the same job title in the same province. We supplement distance with task rarity and the novelty of task pairings. These measures capture advertised combinations rather than completed learning by workers. Tables 6 and 7 test this implication.

Two employers can require the same named skills and nonetheless define different jobs because they combine, sequence and verify those skills differently. The historical-title design focuses on this configuration margin. A high distance can emerge when tasks already common in the labour market are rearranged into a new bundle, so rarity and pair novelty need not be individually significant whenever distance rises. This is particularly relevant to AI-assisted work: drafting, checking, interpreting and communicating an output can remain recognisable activities while their relative weights and joint appearance change. The interpretation remains a hypothesis about employer work design, not a direct measurement of the worker's ability to

integrate them. The component results in Table 6 let the data discriminate between the broader configuration change and a narrower claim that firms predominantly recruit workers for entirely unprecedented tasks.

Retraining and mobility are possible worker responses to a changing task mix, but their payoffs depend on transferability. Evidence using training records for AI-exposed workers^[30] and evaluations of labour-market programmes^[31] highlight this supply-side channel without validating any one training intervention in our data. Because we have no direct observations of training participation, a shift in task distance is evidence about potential demand for integrated abilities, rather than an estimate of a training treatment effect.

Entry screening has competing responses. Complementarity could raise the value of formal education, whereas reduced costs of standardised knowledge tasks could leave the average credential requirement unchanged. AI assistance could narrow the performance gap between novice and experienced workers^[4], yet task integration and checking may strengthen the screening value of experience. We therefore jointly report mean schooling years, degree shares, mean experience years, new-graduate eligibility, and unrestricted requirements. Tasks, skills and wages are linked in earlier work^[32], but the present outcome remains an advertised requirement. Tables 8 and 9 address the third implication.

The entry-margin measures should likewise be read jointly rather than as substitutes. Mean years of schooling captures the central tendency among requirements convertible to years, but bachelor and graduate shares can indicate movement at particular thresholds. An unrestricted-education share records whether a firm states no qualification restriction at all. For experience, the mean among numeric entries can rise while a separate group of postings remains open to recent graduates, or the reverse. These combinations have different implications for access to first jobs. Distinguishing unrestricted from missing text avoids turning an unreported requirement into a firm's explicit decision to waive one. Keeping the schooling and experience denominators separate is essential because a posting may specify only one field. This is why Tables 8 and 9 answer related but not identical empirical questions.

December 2022, the first full month after ChatGPT's public launch, anchors the diffusion period. It is neither the birth of generative AI nor a guarantee that firms had no earlier expectations. Future-value tests, pseudo-treatment dates and pre-trend diagnostics accordingly determine how strongly a discontinuity can be

interpreted. Interactions with services, fiscal self-sufficiency and trade openness describe regional conditions, not separately randomised shocks.

The hypotheses are deliberately conditional rather than directional predictions for all outcomes. If AI substitutes for one task and complements another, the sign of a routine-demand average depends on their relative weights, the response of final demand and employers' willingness to change positions. A task-share movement is predicted more directly by heterogeneous effects, but even there the sign for a particular task depends on how its complements are organised. A distance measure should move if the relative vector of activities changes sufficiently compared with recent postings of the same title; it can remain flat if all tasks are scaled proportionally. Schooling and experience thresholds can move in different directions because credentials and practical knowledge screen different qualities. Stating these conditions in advance matters: it prevents an insignificant average from being re-labelled as a prediction that technology must have no effect.

2.4 A conceptual model and testable margins

Let i , t , k and a denote a province, month, task and capability. Effective demand for task k depends on product demand, the cost-saving substitution margin and augmentation of capabilities used by the task. In equation (1), A_{kt} denotes automation intensity, C_{kt} denotes augmentation intensity, H_{iat} the available capability stock, w_{ka} the task-specific loading and q_{ikt} underlying product demand. The coefficients on automation and augmentation are nonnegative by construction, but their net implication for observed hiring is not predetermined.

$$D_{ikt} = q_{ikt} \exp(-\alpha_k A_{kt} + \gamma_k C_{kt}) \sum_a w_{ka} H_{iat}, \quad \alpha_k, \gamma_k \geq 0 \quad (1)$$

Capability a evolves in equation (2) through pre-entry education E and repeatable post-entry investment U_t , which includes training, practice and transfers across tasks. The depreciation rate is δ_a ; ϕ_a and ψ_a convert the respective investments into future capability. The discounted decision problem in equation (3) equates marginal investment costs to expected returns from future task demand. These expressions organise the hypotheses; vacancy records do not identify the worker-level costs or structural parameters.

$$H_{ia,t+1} = (1 - \delta_a) H_{iat} + \phi_a E_{ia} + \psi_a U_{iat} \quad (2)$$

$$\max_{E, \{U_t\}} \left\{ -c^E(E) + \sum_t \beta^t [R_t(H_t) - c^U(U_t)] \right\} \quad (3)$$

233 Four observable demand margins connect the conceptual mechanism to the empirical tests: persistence of
234 routine tasks, changes in task weights, changes in within-vacancy bundles, and screening by schooling or
235 experience. Figure 1 presents their links to possible incentives for pre-entry accumulation and continuing
236 capability updating. Its arrows denote theoretical relationships rather than estimated worker transitions.

237 The conceptual model also makes clear what cannot be inferred from the reduced-form regressions. The
238 coefficients on technology in equations (1)–(3) are not directly estimated; the data do not supply each
239 employer’s output price, AI capital stock or training technology. A positive experience threshold need not
240 imply that the productivity of experienced workers has risen by the same amount, nor that formal schooling
241 produces less valuable skills. Screening can respond to application volume, uncertainty about quality or
242 changes in the mix of vacancies. Conversely, an unchanging average schooling threshold can conceal higher
243 demand for particular disciplines or credentials, offset by relaxation elsewhere. The empirical implications
244 therefore concern observable employer requirements and their association with exposure, while the investment
245 interpretation describes a plausible incentive that must eventually be tested with worker-level choices and
246 returns.




Figure 1. Analytical framework for generative AI exposure and four margins of skill demand

Note: Solid arrows indicate hypothesised effects on employer requirements; dashed arrows represent possible incentives for education and continuing skill formation. Vacancy data do not measure realised investment.

H1: The average routine-demand response is ambiguous, while task shares can be reallocated across exposure groups (Tables 4–5).

H2: Task-bundle distance may increase even if rarity and pair novelty do not respond uniformly (Tables 6–7).

H3: Schooling and experience thresholds may react differently; neither is a direct measure of realised investment (Tables 8–9).

H4: Diffusion timing and regional economic structure may condition these changes, and dynamic diagnostics must qualify their causal interpretation (Tables 10–13 and Appendix Tables A1–A16).

3 Data, measurement and research design

3.1 Vacancy records, observation units and sample formation

The standardised source contains approximately 5.698 million online postings issued in China between 2018 and 2025. We harmonise publication month, province and posting identifier, separate task descriptions from qualifications, and match job responsibilities to a fixed ontology of 2,083 tasks. The underlying units are postings and posting-task links; empirical units additionally include province-months, province-month-job titles and province-month-tasks. A match row, a posting with a valid matched task, and a regression observation are distinct counts. All tables report their actual sample size rather than treating ontology matches as independent vacancies.

A province-month panel covers 31 provincial units from January 2018 to December 2025 where its respective outcome is defined. The combination design begins in January 2019 because it requires prior postings to construct a twelve-month comparison profile. The education and experience task regressions use different valid sets because a requirement can be intelligible in one field and unavailable in the other. The routine-title analysis likewise requires a valid title, task hit and upstream posting weight. These outcome-specific denominators explain why observation counts differ across Tables 4–9.

The fixed ontology provides a common coordinate system for job descriptions at different dates, mitigating changes in task naming^[29]. We preserve a fixed text-cleaning procedure, task-matching threshold and within-posting deduplication rule. For combination measures, the maximum nonnegative match score is retained for each posting-task link and normalised within the posting. Education and experience at task level instead distribute each posting equally across its valid tasks; routine demand uses separately generated posting weights. A common name for a task does not imply that these distinct weighting rules may be interchanged.

Several practical filters are particularly consequential when reproducing these measures. A single posting can mention a task more than once, and a single description can match several ontology entries. The posting-task link must therefore be deduplicated before any within-posting score is normalised; otherwise repeated wording would mechanically change the task vector. Dates must use the posting month rather than a retrieval or processing date, and a province must be assigned consistently to the recruiting position. A standardised title prevents spelling or formatting variants from generating separate historical comparison groups. For qualification fields, the procedure preserves the original category as well as the numeric mapping so that

288 “unrestricted”, a recognised credential and unavailable text remain distinguishable. These rules underpin
289 comparability but cannot completely remove selection caused by platform coverage or changes in firms’
290 posting practices.

291 The different units also shape how uncertainty should be interpreted. At the posting level, task links from
292 one vacancy share text and are not independent observations of employer decisions. At the province-month
293 level, repeated observations within a province can share economic and platform-specific shocks. At the task
294 level, a national change in the way a task is described can appear simultaneously in many provinces. This
295 motivates the fixed-effects and clustering choices stated below, but it also means that the apparent size of an
296 expanded posting-task file should never be presented as the number of statistically independent treatment
297 assignments. The descriptive and regression tables carry separate N values for exactly this reason.
298 Reproducibility requires preserving the source posting identifier, the mapping from each link to that posting
299 and the specific step at which observations become ineligible for an outcome.

300 The event indicator equals one from December 2022 onward. Regional generative-AI exposure combines
301 a pre-period provincial task structure with a predetermined national diffusion series. Task exposure and
302 regional pre-period exposure are held fixed before treatment in the task-level designs. These constructs
303 measure potential exposure, not direct enterprise adoption, displacement or training^[33,34]. National shocks,
304 composition changes and anticipatory recruitment remain possible limitations; the fixed effects and diagnostic
305 checks below address different parts of this concern.

306 Exposure is fixed with respect to the post-treatment task composition to prevent a mechanical relationship
307 in which changes in the outcome are simultaneously used to redefine the treatment. The regional series
308 combines cross-province baseline variation with a common time-varying diffusion component; the task-level
309 interaction adds heterogeneity among tasks within each province-month. The pre-period anchoring reduces
310 one form of simultaneity, but it does not make baseline task structures randomly assigned. Provinces with
311 different initial task mixes may also differ in unobserved development paths or recruiting platforms. Province
312 effects, month effects and the pairwise fixed effects in task regressions absorb specified shared differences,
313 while pre-trend and placebo exercises check the remaining identifying assumptions. Readers should interpret

the exposure coefficient as a conditional association with a plausibly relevant diffusion index when these assumptions are uncertain.

3.2 Four indicators and their denominators

Write j for a vacancy, c for its standardised title, K for the fixed task set, and r for its nonnegative task-match score. Equation (4) defines the within-posting share q_{jk} used only to construct task-combination measures. The frozen upstream province-month-task demand input is an independent quantity; it must not be inferred from this within-posting normalisation, nor assumed to have the same hiring weights as the routine index.

$$q_{jk} = \frac{r_{jk}^+}{\sum_{\ell \in \mathcal{K}_j} r_{j\ell}^+}, \quad r_{jk}^+ = \max\{r_{jk}, 0\} \quad (4)$$

Routine-task demand. Fix the pre-treatment routine-task identification rule, and let $\mathcal{J}_{itc}$ be the eligible postings for province i , month t and title c . Let ω_j denote the upstream posting weight, H_j the indicator that any routine task was hit, M_j the number of routine hits, and W the sum of eligible weights. Equation (5) forms a weighted routine-posting share and a weighted routine-hit intensity. Their denominators coincide within each title cell, and the sum of posting weights also serves as the regression weight. Routine task content is related to the task-based account of technological change^[16]; it is not a measure of workers' historical spending on schooling.

$$R_{itc}^S = \frac{\sum_{j \in \mathcal{J}_{itc}} \omega_j H_j}{W_{itc}}, \quad R_{itc}^I = \frac{\sum_{j \in \mathcal{J}_{itc}} \omega_j M_j}{W_{itc}} \quad (5)$$

Task reallocation. Let the raw province-month share of a task equal its frozen task-demand count divided by total task demand in that province-month. Equation (6) shrinks it toward the same-month national task share with weight λ_{it} determined by the number of valid provincial postings and the fixed smoothing constant 500. All 2,083 tasks remain on the support, with an unobserved task count treated as zero when the province-month itself exists. We do not impute a missing previous-year comparison. Because shares sum to one, a falling share is a relative change and cannot by itself show falling recruitment in levels^[35].

$$\tilde{S}_{ikt} = \lambda_{it} S_{ikt} + (1 - \lambda_{it}) S_{kt}^N, \quad \lambda_{it} = \frac{N_{it}}{N_{it} + 500} \quad (6)$$

Equation (7) compares a smoothed task share with the corresponding month twelve months earlier and assigns half of the absolute change to the task's reallocation contribution. Equations (8)–(9) separate positive

and negative contributions; their aggregate components are constructed over the same fixed ontology. This decomposition distinguishes redistribution from a parallel scaling of all task counts. The task-exposure quartiles used in the subsequent triple difference are fixed before the diffusion period.

The task-share construction matters for the distinction between an absolute movement and its direction. Half the absolute change assigns a contribution to each task while preventing the aggregate measure from double-counting the same redistribution when gains and losses sum to zero across the complete support. The positive and negative components describe where relative shares expand and contract; they should not be labelled job creation and job destruction. The use of the same calendar month one year earlier also avoids conflating predictable month-of-year differences with reallocation, provided both observations exist. Shrinkage toward the national distribution prevents a small provincial task cell from carrying the same apparent precision as a large one. The constant 500 is a frozen design parameter rather than a coefficient selected after observing the reported regression results.

The interpretation of compositional shares merits particular care when discussing adjustment investment. Suppose a worker's existing knowledge is related to a task whose share falls. That observation alone does not tell us whether demand for the worker falls: related tasks may become more important, and a worker can apply the same underlying ability in another configuration. A negative relative share instead identifies a change in the portfolio of tasks advertised in the provincial market. If employers increase the weight of judgement or verification within a posting, skills accumulated in earlier work may retain value when combined with new practices. The policy implication is therefore more specific than a call for indiscriminate retraining. It concerns the ability to recognise which tasks are changing and to transfer, coordinate or update capabilities in response. Without worker-panel data, we cannot quantify how much new learning any person would need to make that transfer.

$$\Delta_{ikt} = \tilde{S}_{ikt} - \tilde{S}_{ik,t-12}, \quad A_{ikt} = \frac{1}{2} |\Delta_{ikt}| \quad (7)$$

$$A_{ikt}^+ = \max\{\Delta_{ikt}, 0\}, \quad A_{it}^+ = \sum_{k \in \mathcal{K}} A_{ikt}^+ \quad (8)$$

$$A_{ikt}^- = \max\{-\Delta_{ikt}, 0\}, \quad A_{it}^- = \sum_{k \in \mathcal{K}} A_{ikt}^- \quad (9)$$

Within-job combinations. For each vacancy, we compare its q_{jk} vector with the normalised profile of vacancies with the same title in the same province during the preceding twelve months, excluding the current month. Equation (10) is one minus their cosine similarity. A vacancy is excluded from valid distance estimation if the historical title profile has too few eligible postings under the frozen rule or either vector norm is zero. This statistic is a distance between current and historical task vectors, not a newly invented matrix of pairwise task distances.

$$d_j = 1 - \frac{\sum_{k \in \mathcal{K}} q_{jk} p_{i,c(j),t,k}^{(-12,-1)}}{\sqrt{\sum_{k \in \mathcal{K}} q_{jk}^2} \sqrt{\sum_{k \in \mathcal{K}} (p_{i,c(j),t,k}^{(-12,-1)})^2}} \quad (10)$$

The rarity component weights the negative log of a smoothed task frequency in the preceding twelve months by the current vacancy's task shares. Pair novelty measures the fraction of pairs among the prespecified top-weight tasks in a vacancy that did not occur in the provincial twelve-month history. Each component is averaged over the same eligible postings with all three measures defined in a province-month; each is then standardised using its January 2019–November 2022 mean and sample standard deviation. Their equally weighted average is the composite combination measure. The descriptive combination statistic in Table 1 retains its unstandardised historical distance definition and therefore has a different mean from the standardised regression outcomes in Table 6.

The formal combination window offers $31 \times 84 = 2,604$ potential province-month cells from 2019 to 2025. Twelve are absent from the parent empirical panel, concentrated in October and November 2021. Of the remaining 2,592 cells, 1,783 have identifiable combination measures and enter the formal regression; 809 do not. Thus the regression covers 68.5% of all potential cells. Annual shares of potential cells represented are 72.8%, 56.5%, 67.2%, 82.8%, 58.9%, 60.2% and 80.9% for 2019–2025. We neither code missing combinations as zero nor impute them. Available province-month summaries do not permit a mutually exclusive attribution of the 809 missing outcome cells to upstream matching stages, so the estimates pertain to the identifiable sample.

This coverage calculation is more informative than reporting only a complete-case regression count. Its numerator is the 1,783 province-month cells with an identifiable outcome; its reference set is all 2,604

province-month opportunities in the seven-year formal window. The 12 cells absent even from the parent panel and 809 cells without a valid combination measure are conceptually different exclusions. The latter can reflect requirements of the historical same-title comparison, vector validity and simultaneous availability of the three components, but available summaries cannot uniquely partition them among these channels. The marked differences in annual coverage mean that an aggregate coefficient need not describe a balanced provincial population at every date. Minimum-posting thresholds and province omission check specific vulnerabilities in the observed cells; neither reconstructs the missing counterfactual combinations. This limitation is especially pertinent when giving the coefficient an economy-wide interpretation.

Entry requirements. Valid education categories map to completed schooling years as follows: lower secondary or less 9; upper secondary and vocational secondary 12; college 15; bachelor 16; master or unspecified graduate 19; doctorate 22. Unrestricted education and genuinely missing education are separate statuses, neither assigned a schooling year. New graduates, students and internships map to zero required years of prior experience; unrestricted and missing experience do not enter the mean-years denominator. As in equation (11), each posting contributes equally to its K_j valid tasks, and education and experience have distinct valid posting-task sets and denominators.

$$\begin{aligned} Y_{ikt}^E &= \frac{\sum_{j \in \mathcal{J}_{ikt}^E} a_{jk} e_j}{\sum_{j \in \mathcal{J}_{ikt}^E} a_{jk}}, \\ Y_{ikt}^X &= \frac{\sum_{j \in \mathcal{J}_{ikt}^X} a_{jk} x_j}{\sum_{j \in \mathcal{J}_{ikt}^X} a_{jk}}, \quad a_{jk} = \frac{1}{K_j}. \end{aligned} \tag{11}$$

Equations (12)–(14) define the province-month degree and unrestricted shares, the new-graduate and unrestricted-experience shares, and the two conditional means. Categorical shares use postings for which the relevant education or experience field has an interpretable category; conditional year means use the smaller set with a convertible numeric year value. Hence “unrestricted” contributes to an unrestricted-category share but is never assigned zero years; a new graduate contributes a zero to the valid experience-years mean and is separately identified by the new-graduate share. Missing text is not treated as unrestricted.

$$(B_{it}, M_{it}, U_{it}^E) = \frac{1}{N_{it}^E} \sum_{j \in \mathcal{J}_{it}^E} (\mathbf{1}\{e_j \geq 16\}, \mathbf{1}\{e_j \geq 19\}, \mathbf{1}\{e_j = \text{unrestricted}\}) \tag{12}$$

$$(F_{it}, U_{it}^X) = \frac{1}{N_{it}^X} \sum_{j \in J_{it}^X} (\mathbf{1}\{x_j = \text{new graduate}\}, \mathbf{1}\{x_j = \text{unrestricted}\}) \quad (13)$$

$$\bar{e}_{it} = \frac{1}{N_{it}^{E, \text{yr}}} \sum_{j \in J_{it}^{E, \text{yr}}} e_j, \quad \bar{x}_{it} = \frac{1}{N_{it}^{X, \text{yr}}} \sum_{j \in J_{it}^{X, \text{yr}}} x_j \quad (14)$$

3.3 Exposure, estimation and statistical inference

Baseline province-month specifications use equation (15), where Y_{it} is a defined province-month outcome, G_{it} is regional generative-AI exposure, $X_{it}^\top$ includes only the controls actually entered in the corresponding reported specification, and μ_i and τ_t are province and month fixed effects. Potential economic, service-sector, fiscal and trade covariates are included where estimable; specification labels in tables report their actual inclusion. Province-by-month fixed effects in task regressions absorb common province-month controls, so those controls cannot be independently identified alongside that fixed effect. The routine-title outcome instead uses title-cell weights, province-by-title fixed effects and month effects, and should not be redescribed as equation (15) estimated on province-months.

The fixed-effect structure must also be matched to the observational unit. Province and month indicators in a province-month regression account for time-invariant provincial differences and shocks common to all provinces in a month. A province-by-month effect in the task-level equation is much more demanding: it compares tasks within a given province at a given date and thereby absorbs any variable shared by all those tasks in that province-month, including the main effect of a provincial monthly covariate. Task-by-month effects account for national changes specific to a task, and province-by-task effects absorb persistent local specialisation. The remaining identifying interaction must vary jointly across the three dimensions. An automatically generated table row that merely says “fixed effects: yes” would obscure these differences; the English tables retain their actual listed effects.

$$Y_{it} = \beta G_{it} + X_{it}^\top \theta + \mu_i + \tau_t + \varepsilon_{it} \quad (15)$$

For task reallocation, equation (16) compares each fixed task-exposure quartile from two to four with the first quartile under regional pre-period exposure and the post indicator. It includes province-by-task, province-by-month and task-by-month fixed effects. For education and experience, equation (17) replaces quartile indicators with fixed continuous task exposure under the same three dimensions of fixed effects. Identification thus comes from the interaction of within-province-month task exposure with provincial baseline exposure

after diffusion, net of the stated pairwise effects. Province and task are separate clustering dimensions for the task-level estimates.

$$Y_{ikt} = \sum_{g=2}^4 \beta_g I_k^{(g)} Z_i P_t + X_{ikt}^T \theta + \alpha_{ik} + \delta_{it} + \gamma_{kt} + \varepsilon_{ikt} \quad (16)$$

$$Y_{ikt}^h = \beta_h E_k Z_i P_t + \alpha_{ik} + \delta_{it} + \gamma_{kt} + \varepsilon_{ikt}^h, \quad h \in \{E, X\} \quad (17)$$

Routine-demand models cluster at province level. Combination models absorb province and calendar-month effects and cluster by province, with a wild-cluster bootstrap as a sensitivity check when the number of provincial clusters is limited. Task-level models cluster by both province and task. Some placebo dates, leads and longer pre-treatment windows reveal anticipatory movements. A lead test diagnoses pre-existing adjustment rather than establishing that the post coefficient is necessarily biased in a known direction; leave-one-province-out checks diagnose leverage, and alternative exposure measures assess measurement sensitivity. These checks cannot be collapsed into a claim that every form of endogeneity is eliminated. The empirical novelty lies primarily in the consistent task ontology, frozen denominators, historical title profiles and decomposition of demand contributions, not in a new regression estimator.

Statistical inference is another reason the four sets of stars cannot be compared mechanically. Province clustering permits arbitrary dependence among observations from the same province under the conventional procedure, whereas two-way clustering in the task regressions additionally addresses shared task shocks. The combination design has a limited number of provincial clusters, making wild-cluster resampling a useful sensitivity analysis for its key coefficients. A wild p-value can differ materially from a conventional p-value even though the coefficient and sample have not changed. We consequently privilege the distance estimate that retains support under the bootstrap and label the composite as less secure. For other outcomes, we do not impute unreported bootstrap significance or treat a conventional star as proof of robust rejection under every dependence structure.

An AI-assisted writing tool was used to help revise English phrasing and table formatting. The author remains responsible for the analysis and the final manuscript.

4 Empirical results

4.1 Distributions and correlations

Table 1 displays the distributions of the demand margins, entry requirements and principal controls. The routine, reallocation and combination variables measure different objects on different scales, so their numerical means cannot be compared as intensities of one latent investment. The average routine measure is approximately 0.0382, the reallocation measure 0.1670 and the unstandardised combination descriptive measure 0.5335. Mean advertised schooling is about 15.33 years, approximately 53.64% of postings require at least a bachelor's degree, mean required experience is approximately 2.44 years, and about 21.32% state no experience restriction. These distributions motivate examining both mean thresholds and unrestricted categories.

The descriptive sample sizes also contain substantive information. Postings and province-months differ because one cell can represent many vacancies; title-level routine regressions can contain far more rows than province-month models; task-level designs can contain millions of province-task-month observations derived from the task structure. Large numbers of linked task rows do not increase the number of independent provinces. The descriptive combination series and the regression combination series also differ in standardisation and validity criteria. Reading the reported N together with the measure definition guards against an otherwise tempting but invalid comparison of a descriptive mean with a coefficient from a different sample. We use Table 1 to check the distribution of each signal before moving to models that condition on fixed effects and control sets.

Correlation tables are included to document the empirical setting, not to substitute for the research design. A pairwise relationship mixes permanent differences across provinces, national trends and the post-period variation used in the main models. It can even have a different sign from a conditional estimate after fixed effects change the comparison set. By displaying two panels of correlations we allow readers to inspect the association among demand indicators and macroeconomic controls without squeezing a seventeen-variable matrix into illegible type. The formal conclusions draw on the model for each outcome, its valid observations and the corresponding clustered uncertainty. Where a coefficient is small, we state its original-scale meaning alongside any display multiplier so that the result is not mistaken for an effect of the multiplier itself.

Table 1. Descriptive statistics

Variable	N	Mean	SD	Minimum	P25	Median	P75	Maximum
Routine-task demand	2964	0.038	0.024	0.000	0.027	0.035	0.046	0.321

Task-reallocation demand	2580	0.167	0.074	0.046	0.111	0.144	0.216	0.586
Task-combination demand	2916	0.533	0.045	0.194	0.516	0.535	0.551	0.880
Required schooling (years)	2964	15.330	0.393	12.854	15.169	15.350	15.525	19.000
Bachelor's degree or above	2964	0.536	0.129	0.000	0.453	0.530	0.612	1.000
Graduate degree or above	2964	0.030	0.040	0.000	0.012	0.024	0.038	1.000
Share with no education restriction	2964	0.095	0.086	0.000	0.038	0.075	0.132	0.774
Required experience (years)	2962	2.438	0.448	0.000	2.200	2.460	2.715	5.500
Fresh graduate share	2964	0.065	0.077	0.000	0.017	0.037	0.094	1.000
Share with no experience restriction	2964	0.213	0.135	0.000	0.127	0.207	0.276	1.000
Regional generative AI exposure shock	2964	0.000	1.185	-9.512	0.000	0.000	0.000	6.978
Log real GDP per capita	2914	8.690	0.415	7.734	8.396	8.618	8.923	9.923
Fiscal self-sufficiency	2914	0.486	0.310	0.023	0.285	0.415	0.600	3.606
Trade openness	2914	23.555	21.700	0.136	8.823	14.962	31.909	110.235
Child dependency ratio	2542	24.182	6.437	11.610	18.950	24.990	28.409	38.380
Old-age dependency ratio	2542	19.424	5.087	8.040	15.120	19.790	23.030	32.090
Service-sector share	2542	0.558	0.091	0.408	0.513	0.540	0.573	0.975

Notes: N denotes valid observations. Proportion measures use their original units. The descriptive task-combination measure is the unstandardised distance based on 2916 province-month observations; the regression outcomes in Table 6 use pre-treatment standardisation and therefore have different distributions.

Tables 2 and 3 divide the pairwise correlation matrix into readable panels. Simple correlations of the exposure measure with combination demand and mean education and experience are positive but generally modest. Controls such as real income, openness and the service-sector share show some correlation, yet pairwise correlations alone neither establish nor refute harmful multicollinearity in the fixed-effects specifications. Degree shares also correlate because they represent nested qualifications; we do not enter them together as competing explanatory variables in one main equation.

Table 2. Correlations among primary measures

Variable	1	2	3	4	5	6	7	8	9
1	1.000								
2	0.029	1.000							
3	-0.068***	-0.129***	1.000						
4	0.246***	0.018	-0.105***	1.000					
5	0.264***	0.013	-0.032*	0.775***	1.000				
6	0.165***	0.059***	-0.015	0.552***	0.438***	1.000			
7	-0.140***	0.092***	0.108***	-0.322***	-0.203***	-0.171***	1.000		
8	0.059***	-0.114***	-0.152***	0.289***	0.315***	0.021	-0.130***	1.000	
9	0.063***	0.147***	-0.096***	0.071***	-0.013	0.254***	-0.199***	-0.383***	1.000
10	-0.160***	-0.012	0.164***	-0.325***	-0.199***	-0.155***	0.691***	-0.128***	-0.448***
11	0.005	-0.038*	0.049***	0.049***	0.079***	0.045**	-0.010	0.080***	-0.003
12	0.242***	-0.039**	-0.006	0.483***	0.530***	0.486***	-0.321***	0.278***	0.078***
13	0.110***	-0.091***	0.054***	0.226***	0.328***	0.253***	-0.032*	0.164***	-0.025
14	0.165***	-0.103***	0.060***	0.375***	0.488***	0.394***	-0.092***	0.240***	-0.039**
15	-0.121***	-0.009	-0.052***	-0.274***	-0.342***	-0.330***	0.068***	-0.051**	-0.008
16	0.108***	0.102***	-0.066***	0.284***	0.273***	0.331***	-0.259***	0.127***	0.070***
17	0.210***	-0.065***	-0.027	0.378***	0.457***	0.416***	-0.105***	0.241***	-0.080***


Notes: Variable numbers are (1) Routine-task demand; (2) Task-reallocation demand; (3) Task-combination demand; (4) Required schooling (years); (5) Bachelor's degree or above; (6) Graduate degree or above; (7) Share with no education restriction; (8) Required experience (years); (9) Fresh graduate share; (10) Share with no experience restriction; (11) Regional generative AI exposure shock; (12) Log real GDP per capita; (13) Fiscal self-sufficiency; (14) Trade openness; (15) Child

507 dependency ratio; (16) Old-age dependency ratio; (17) Service-sector share. Table 2 displays rows
508 (1)–(17) against columns (1)–(9); Table 3 displays rows and columns (10)–(17). Correlations are
509 rounded to three decimal places.
510 $*p < .10$, $**p < .05$, $***p < .01$

511 **Table 3. Correlations among additional measures**

Variable	10	11	12	13	14	15	16	17
10	1.000							
11	-0.028	1.000						
12	-0.311***	0.210***	1.000					
13	-0.057***	0.077***	0.499***	1.000				
14	-0.111***	0.227***	0.749***	0.588***	1.000			
15	0.042**	-0.054***	-0.482***	-0.332***	-0.502***	1.000		
16	-0.200***	-0.196***	0.306***	0.162***	0.187***	-0.417***	1.000	
17	-0.127***	0.151***	0.564***	0.369***	0.691***	-0.353***	0.106***	1.000

512 Notes: Variable numbers (10)–(17) follow the complete variable key beneath Table 2. This panel
513 completes the correlation matrix. Coefficients and significance symbols use the same conventions as
514 Table 2.



515

516 **Figure 2. Main regression coefficients for the four margins of skill demand**

517 Note: Dots represent reported main coefficients and bars conventional approximate 95% intervals.
518 Each panel has its own horizontal scale. Robust inference for task combinations appears in Figure 3
519 and Table 7.

520 4.2 Routine demand as a persistence benchmark

521 Table 4 estimates both the weighted routine-posting share and weighted hit intensity at the province-
522 month-title level. A universal displacement story would predict declining routine recruitment, while scale and
523 augmentation could offset that decline. The estimates are approximately 0.000037 and 0.000134 on the
524 original outcome scale, respectively; neither reaches conventional significance in the reported main

specification. The size and precision therefore do not establish an economy-wide increase or decline in these routine requirements. This is a persistence benchmark, rather than proof that every routine task is unaffected at every date.

Table 4. Routine-task demand

Variable	Routine-demand share (scaled)	Routine-demand intensity (scaled)
Regional generative AI exposure shock	0.371 (2.681)	1.339 (3.000)
Controls	No	No
Fixed effects	province $\times$ job title, month	province $\times$ job title, month
N	2184097	2184097
R^2	0.431	0.433

Notes: Standard errors clustered by province are in parentheses. Coefficients and standard errors use the scaling stated in the column headings; observations and R^2 are unscaled.
 $*p < .10$, $**p < .05$, $***p < .01$

Alternative exposure constructions, leads and province omission appear in Appendix Tables A2–A3 and A6. Some dynamic checks are less reassuring over longer windows. A nonsignificant coefficient is also compatible with effects within its confidence interval, so the paper does not infer that a perfectly zero routine effect has been identified. Its more specific empirical question is whether other task dimensions change even when the routine average does not move reliably.

4.3 Reallocation of task shares

Table 5 separates absolute task reallocation into expanding and contracting task-share contributions across the three higher exposure quartiles relative to the first. The interaction of pre-period task exposure, baseline regional exposure and the post period creates differences in changes among tasks within a common province-month. The results vary across exposure groups; expansion and contraction are not simply mirror images in the reported contribution outcomes. In the frozen task-share design, highly exposed tasks experience a lower relative share, pointing to a redistribution of requirements across tasks rather than a proportional change in total recruiting.

In interpreting this decomposition, the quartile coefficients describe differences relative to the omitted first exposure group after the interacted regional and time dimensions are accounted for. They are not estimates of the absolute number of positions gained by each quartile. A positive contribution to absolute reallocation signals a stronger movement of shares, whether the underlying direction is gain or loss; the expansion and contraction columns clarify that direction. Even those directional columns concern fractions of

demand over the fixed task set, not employment separations or wages. The empirical distinction helps reconcile an imprecise average routine coefficient with appreciable changes in the composition of tasks sought by employers. It does not establish that every individual worker must change occupations or make a new training purchase.

Table 5. Task-reallocation components

Variable	Absolute reallocation contribution (scaled)	Expansion contribution (scaled)	Contraction contribution (scaled)
Exposure group 2 × regional exposure × post	2.662** (1.003)	2.672*** (0.783)	2.653* (1.347)
Exposure group 3 × regional exposure × post	2.026*** (0.620)	2.420*** (0.601)	1.633* (0.848)
Exposure group 4 × regional exposure × post	0.598 (0.713)	1.084 (0.659)	0.113 (1.001)
Controls	Absorbed	Absorbed	Absorbed
Fixed effects	province × task; province × month; task × month	province × task; province × month; task × month	province × task; province × month; task × month
N	5374140	5374140	5374140
R ²	0.860	0.852	0.831

Notes: Standard errors are clustered by province and task. Coefficients and standard errors use the scaling specified in the headings; observations and R² are unscaled. Province-by-month fixed effects absorb common province-month covariates. Significance symbols refer to unrounded estimates.

p < .10, **p < .05, *p < .01*

A change in relative task weights creates a potential reason for incumbents to transfer skills or adapt job practices. It does not reveal whether a worker actually enrolled in training, nor whether the causal path runs entirely through the December 2022 event. Appendix Table A1 details stage effects, while Appendix Table A7 includes diagnostic specifications. Some future-value and reverse-prediction results suggest adjustment may have begun earlier. Accordingly, the evidence supports task reconfiguration during diffusion with a clearly stated anticipation boundary.

4.4 Within-job task combinations

Task reallocation concerns how demand is distributed across the ontology; the combination outcomes ask whether tasks appear together in different ways inside postings. Table 6 reports the equal-weight composite, current-versus-history cosine distance, rarity, pair novelty and an alternative distance. The main distance coefficient is 0.078683 and the composite coefficient is 0.047027, both positive under conventional province-clustered inference; the alternative distance is also positive. Rarity and pair novelty do not display similarly stable significance. This pattern places the strongest weight on a change in the overall task vector rather than a blanket increase in unusual individual tasks.

Table 6. Within-job task combinations

Variable	Composite combination index	Task-combination distance	Task rarity	Task pairing	Alternative distance (scaled)
Regional generative AI exposure shock	0.047** (0.022)	0.079*** (0.028)	0.042 (0.037)	0.020 (0.025)	2.081*** (0.535)
Controls	No	No	No	No	No
Fixed effects	province, month	province, month	province, month	province, month	province, month
N	1783	1783	1783	1783	1783
R ²	0.330	0.215	0.550	0.448	0.364

*Notes: Province-clustered standard errors are in parentheses. Only the alternative-distance column uses the indicated scaling; all other columns, observations and R² remain in their original units. Significance symbols refer to conventional clustered inference; see Table 7 for wild-cluster inference. * $p < .10$, ** $p < .05$, *** $p < .01$*

Table 7 assesses inference with a province-level wild-cluster bootstrap. For the main distance measure the bootstrap p-value is approximately 0.025 and the corresponding 95% interval remains above zero. The composite is less secure at the 5% level under this procedure, explaining why distance is the headline combination measure. Appendix Tables A4–A5 and A8 examine trends, alternative exposure and identification. The direction is broadly robust across several specifications, although long-window pre-trends and some pseudo dates limit an unqualified event-based interpretation.

The distance outcome is particularly sensitive to its reference profile. Comparing each posting with the preceding twelve months of the same standardised title within its province reduces cross-occupation compositional variation and asks whether the job's task configuration departs from recent local practice. The title may still contain positions of differing seniority and firms, and changes in the composition of firms posting vacancies can shift the observed profile. Moreover, a large distance is not automatically an improvement in productivity or job quality. The bootstrap and alternative-distance results speak to the sign and inferential stability of an employer-demand measure; the trend and placebo evidence speaks to timing; neither independently verifies an underlying worker investment mechanism. This division of evidence is why the headline emphasises a change in advertised task combinations.

Table 7. Inference for task-combination measures

Measure	Main coefficient	Conventional p-value	Wild-cluster p-value	Wild 95% CI	N
Composite combination index	0.047 (0.022)	0.040	0.115	[-0.027, 0.100]	1783
Task-combination distance	0.079 (0.028)	0.008	0.025	[0.018, 0.153]	1783
Task rarity	0.042 (0.037)	0.259	0.416	[-0.084, 0.105]	1783
Task pairing	0.020 (0.025)	0.417	0.415	[-0.028, 0.096]	1783

Alternative distance	0.002 (0.001)	<0.001	0.014	[0.001, 0.003]	1783
----------------------	------------------	--------	-------	----------------	------

Notes: Each measure has a coefficient and, on the following row, its province-clustered conventional standard error in parentheses. Wild-cluster bootstrap p-values and 95% confidence intervals are reported separately. Table 7 summarises inference for models in Table 6; their R^2 values are reported in Table 6 and are not repeated here. Very small nonzero values use scientific notation. $*p < .10$, $**p < .05$, $***p < .01$




Figure 3. Coefficients and robust confidence intervals for task-combination demand

Note: Intervals are wild-cluster bootstrap 95% confidence intervals from Table 7. The alternative distance is displayed on the scale indicated in its panel; coefficients on different axes should not be compared by size.

Appendix Table A24 changes the minimum number of identifiable postings per province-month to five and ten, using the same province and month effects, province clustering and 9,999 wild-cluster replications. Distance coefficients across the original, five-posting and ten-posting samples are 0.0787, 0.0775 and 0.0746; their wild p-values are 0.0252, 0.0193 and 0.0020, with intervals above zero. Composite p-values are 0.1146, 0.0704 and 0.0113, respectively. Increasing the threshold reduces observations from 1,783 to 1,293 and provinces from 31 to 28. The distance association does not depend on low-posting cells within the identifiable samples, but changed geographic coverage prevents a general conclusion about missing cells.

4.5 Education and prior-experience requirements

Table 8 compares advertised mean years of education and experience in task-level triple differences with province-by-month, province-by-task and task-by-month effects. The education coefficient is 0.004770 with a standard error of 0.004404 and is not conventionally significant. The experience coefficient is 0.008334 with a standard error of 0.002094 and is positive and statistically distinguishable from zero in the reported

specification. This is evidence of an increased experience screening requirement in exposed task cells, while the same design does not show a systematic increase in average educational years.

Table 8. Education and experience requirements at the task level

Variable	Required schooling (years)	Required experience (years)
Task exposure × regional exposure × post	0.005 (0.004)	0.008*** (0.002)
Controls	Absorbed	Absorbed
Fixed effects	province × month; province × task; task × month	province × month; province × task; task × month
N	4623068	4451852
R ²	0.735	0.622

Notes: Standard errors are two-way clustered by province and task. Province-by-month fixed effects absorb common contemporaneous province controls.

p < .10, **p < .05, *p < .01*

The two regressions have different valid sets and denominators. Without a common-sample test of the difference between their coefficients, one significant and one nonsignificant estimate do not by themselves establish a statistically significant contrast. Nor does the schooling result imply formal education has no value. Experience may proxy for contextual judgement, integration or verification, but none is independently identified by the years requirement. Lead diagnostics in Appendix Table A10 suggest possible anticipatory experience changes; the headline should remain a recruitment-threshold finding rather than a sharp single-date causal claim.

One should also distinguish a within-task-cell conditional contrast from an economy-wide aggregate. The three-dimensional fixed effects in Table 8 remove common provincial monthly shocks, national task-month developments and persistent province-task differences. The task-exposure interaction then asks whether requirements for relatively exposed tasks change differently in provinces with different baseline exposure. If the distribution of valid task cells changes, however, the target population can change as well. Table 9 instead aggregates postings to province-month outcomes and varies the control set; it gives useful information about the distribution of screening categories but has a different source of identifying variation. A formal comparison of education and experience coefficients would require common units, a shared valid sample and an explicit estimated covariance, none of which is provided by merely juxtaposing the two significance levels.

Table 9 decomposes entry screening into seven outcomes: mean education years, bachelor-or-higher, graduate-or-higher, unrestricted education, mean experience years, new-graduate eligibility and unrestricted experience. Baseline, core, basic and extended control variants are shown as actually estimated. The province-

month measures do not move uniformly; mean schooling is negative in some sparse specifications and loses significance in others, but valid samples also change with controls. This table documents specification sensitivity and category composition. The task-level Table 8 provides the principal within-province-month identification, while Table 9 is descriptive corroboration and decomposition.

Table 9. Entry requirements across specifications

Outcome	Baseline	Core controls	Basic controls	Extended controls
Required schooling (years)	-0.015** (0.007)	-0.014* (0.007)	-0.015 (0.010)	-0.015 (0.010)
N	2964	2914	2542	2542
R ²	0.601	0.641	0.642	0.642
Bachelor's degree or above	-0.001 (0.003)	-0.001 (0.003)	0.005 (0.004)	0.004 (0.004)
N	2964	2914	2542	2542
R ²	0.564	0.611	0.618	0.619
Graduate degree or above	3.150 × 10 ⁻⁴ (0.001)	0.001 (0.001)	-4.360 × 10 ⁻⁴ (0.001)	-0.001 (0.001)
N	2964	2914	2542	2542
R ²	0.286	0.479	0.441	0.443
Share with no education restriction	0.002 (0.002)	0.003* (0.002)	0.001 (0.004)	0.001 (0.004)
N	2964	2914	2542	2542
R ²	0.638	0.639	0.622	0.622
Required experience (years)	0.006 (0.013)	0.005 (0.012)	0.012 (0.018)	0.011 (0.018)
N	2962	2914	2542	2542
R ²	0.454	0.496	0.515	0.515
Fresh graduate share	0.001 (0.002)	0.002 (0.002)	0.002 (0.004)	0.002 (0.004)
N	2964	2914	2542	2542
R ²	0.543	0.627	0.625	0.626
Share with no experience restriction	0.002 (0.003)	0.002 (0.003)	0.001 (0.005)	0.001 (0.005)
N	2964	2914	2542	2542
R ²	0.691	0.728	0.702	0.703
Controls	No	Core controls	Basic controls	Extended controls
Fixed effects	province, month	province, month	province, month	province, month

Notes: Outcome-specific N and R² follow Tables A17–A23. Parentheses: province-clustered standard errors.

*p < .10, **p < .05, ***p < .01

5 Heterogeneity and identification diagnostics

5.1 Province-month heterogeneity in aggregate task reallocation

The supplementary heterogeneity regressions in Tables 10–12 use a province-month panel with six monthly controls, province and month fixed effects, and province-clustered standard errors. They do not reuse the estimation units or samples of all three corresponding main tables. In Table 10, the dependent variable is the province-month aggregate reallocation measure (2,158 observations), whereas Table 5 estimates task-level triple differences on 5,374,140 province-month-task observations. Table 10 finds positive interactions with pre-treatment service-sector share, fiscal self-sufficiency and trade openness. These are conditional

associations for the aggregate outcome; they are not heterogeneity estimates for the Table 5 task-level coefficients.

Table 10. Province-month heterogeneity in aggregate task reallocation

Variable	Service-sector share	Fiscal self-sufficiency	Trade openness
Regional generative AI exposure shock	0.005*** (0.001)	0.005*** (0.001)	0.005*** (0.001)
Regional generative AI exposure × service sector	0.002*** (3.210×10^{-4})		
Regional generative AI exposure × fiscal self-sufficiency		0.002403*** (0.000462)	
Regional generative AI exposure × trade openness			0.002851*** (0.000609)
Controls	Yes	Yes	Yes
Fixed effects	province, month	province, month	province, month
N	2158	2158	2158
R ²	0.890	0.891	0.891

Notes: The dependent variable is the province-month aggregate reallocation measure. Each moderator is a standardised provincial mean measured through October 2022 and interacted with regional generative AI exposure. All columns include six monthly controls, province and month fixed effects; standard errors are clustered by province. The sample contains 2158 province-month observations. Table 5 instead uses province-month-task cells and a task-level triple-difference design; its coefficients are not directly comparable.

* $p < .10$, ** $p < .05$, *** $p < .01$

5.2 Province-month combination and routine-demand outcomes

Table 11 uses the earlier task-combination distance measure in a province-month control-complete sample (2,505 observations). In the available frozen P2 panel, the earlier aggregate combination-distance variable agrees exactly with the alternative distance measure reported in column (5) of Table 6 wherever both are observed. The headline P2 distance in Table 6, column (2), is a different outcome; its formal P2 sample has 1,783 observations. Table 11 therefore offers contextual heterogeneity evidence for the earlier distance measure, not a test of heterogeneity in the headline P2 distance. Among the three interactions, only fiscal self-sufficiency is significant at the conventional 5% level in the reported Table 11 results.

Table 11. Province-month heterogeneity in the earlier combination-distance measure

Variable	Service-sector share	Fiscal self-sufficiency	Trade openness
Regional generative AI exposure shock	0.002 (0.001)	0.002 (0.001)	0.002 (0.001)
Regional generative AI exposure × service sector	1.000×10^{-6} (4.730×10^{-4})		
Regional generative AI exposure × fiscal self-sufficiency		−0.001871** (0.000782)	
Regional generative AI exposure × trade openness			−0.001033 (0.000647)
Controls	Yes	Yes	Yes
Fixed effects	province, month	province, month	province, month
N	2505	2505	2505
R ²	0.234	0.235	0.234

Notes: The outcome is the earlier combination-distance measure, corresponding to the alternative-distance outcome in Table 6, column (5), in the available frozen P2 panel; it is not the headline P2 distance in Table 6, column (2). This table uses 2505 province-month observations with six monthly controls, province and month fixed effects, and province-clustered standard errors. Table 6 restricts all five outcomes to 1783 formal P2 province-month observations and omits these controls. The coefficients and sample sizes must not be compared as a matched-sample heterogeneity test.

Moderators are standardised pre-treatment provincial means.

* $p < .10$, ** $p < .05$, *** $p < .01$

Table 12 examines the province-month aggregate routine-demand share on 2,542 observations. Its

structural interactions are imprecise. Table 4, by contrast, estimates weighted province-month-title regressions on 2,184,097 title cells, absorbing province-by-title and month fixed effects. Thus Table 12 addresses variation in an aggregated outcome across provincial characteristics, rather than heterogeneity in the Table 4 title-level coefficient. Failure to reject the interactions does not establish the absence of heterogeneity.

Table 12. Province-month heterogeneity in aggregate routine-task share

Variable	Service-sector share	Fiscal self-sufficiency	Trade openness
Regional generative AI exposure shock	5.800×10^{-5} (0.001)	1.050×10^{-4} (0.001)	2.030×10^{-4} (0.001)
Regional generative AI exposure $\times$ service sector	0.001 (0.001)		
Regional generative AI exposure $\times$ fiscal self-sufficiency		-0.000572 (0.000978)	
Regional generative AI exposure $\times$ trade openness			0.000051 (0.000916)
Controls	Yes	Yes	Yes
Fixed effects	province, month	province, month	province, month
N	2542	2542	2542
R^2	0.217	0.217	0.216

Notes: The outcome is the province-month aggregate routine-task share; the 2542 observations are province-month cells with complete controls. All models include six monthly controls, province and month fixed effects, and province-clustered standard errors. Table 4 instead uses weighted province-month-title cells, province-by-title and month fixed effects; it also reports routine-task hit intensity as a separate outcome. Accordingly, Table 12 is a supplementary aggregate analysis, not a within-sample interaction test of Table 4. Moderators are standardised pre-treatment provincial means.

* $p < .10$, ** $p < .05$, *** $p < .01$

5.3 Experience screening and dynamic boundaries

For task-level experience thresholds, Table 13 shows a negative service-sector interaction, whereas the fiscal and openness interactions do not establish an equally stable pattern. One possible explanation is that firms in service-intensive regions can pair novice hires with AI assistance for selected assignments, reducing reliance on prior years of experience; this mechanism is compatible with the estimate but is not directly observed. The table establishes conditional variation in an advertised threshold, not the employment outcome of novice hires.

These province-month interactions describe conditional variation in aggregate recruitment measures. The three aggregate outcomes, their complete-case samples and their controls differ from the task-level and title-level principal regressions. Their signs and precision therefore cannot be used to infer that the main task-level reallocation effect, the P2 headline combination-distance effect or the title-level routine effect varies by the same provincial characteristic. A matched-sample interaction regression on each corresponding principal outcome would be required for those claims.

Table 13. Heterogeneity for experience requirements

Variable	Service-sector share	Fiscal self-sufficiency	Trade openness
Triple-difference effect	0.009*** (0.002)	0.010*** (0.002)	0.009*** (0.002)
Triple-difference effect × service-sector share	-0.002** (0.001)		
Triple-difference effect × fiscal self-sufficiency		-0.002 (0.002)	
Triple-difference effect × trade openness			-0.001 (0.002)
Controls	Absorbed	Absorbed	Absorbed
Fixed effects	province × task; province × month; task × month	province × task; province × month; task × month	province × task; province × month; task × month
N	4451852	4451852	4451852
R ²	0.622	0.622	0.622

Notes: Standard errors are two-way clustered by province and task. Province-by-month fixed effects absorb common contemporaneous province controls.

p < .10, **p < .05, *p < .01*

Figure 4 displays stage-specific task-reallocation coefficients across exposure groups. The plot is based on the reported staged triple differences and conventional clustered intervals; it is not a fully saturated monthly event study. Appendix Tables A1–A16 include lead values, reverse prediction, alternative capability-frontier Bartik exposure, leave-one-province-out exercises, placebo timing and regional splits. Each addresses a distinct threat: anticipation, treatment measurement, dependence on a province or dynamic specification. A statistically significant alternative exposure result cannot repair a failed pre-trend in a different outcome.

These diagnostics also prevent a binary “passes all checks” narrative. A future-value coefficient can be consistent with firms anticipating a technological transition, but could also point to time-varying confounding under the proposed specification. An alternative frontier-based exposure measure tests whether one specific index drives the result, not whether all variation is as good as random. Dropping a province assesses undue leverage, while a pseudo-treatment date probes whether similar-looking movements occur outside the chosen

event date. Longer pre-event windows sometimes ask a harder question than a short lead: they may detect gradual changes already present before public release. We disclose these distinctions for the routine, reallocation, combination and entry-threshold outcomes separately, and reserve stronger causal language for designs whose timing and comparison assumptions warrant it.




Figure 4. Stage-specific dynamic coefficients for task-reallocation demand
Note: Points denote stage-specific triple-difference estimates relative to the first task-exposure group; bars are approximate 95% intervals from conventionally clustered standard errors. These estimates are not a monthly event study.

Taken together, the designs provide evidence of changes with varying identification strength. Routine demand lacks a stable average response. Reallocation shows differences among exposure groups but some evidence of early adjustment. Distance in within-job combinations remains positive under province-level wild-cluster inference, subject to incomplete outcome coverage and pre-trend caveats. Mean education does not exhibit a stable task-level increase, whereas experience requirements rise and may also adjust ahead of the public-release date. Reporting these differences avoids forcing four outcomes into one precisely timed causal effect.

6 Conclusions and implications

Matching 2018–2025 Chinese vacancies to a fixed task ontology reveals that the recruitment consequences of generative AI diffusion differ across four demand margins. The main routine specifications show no robust average movement; tasks are redistributed across exposure groups; the distance of new within-job task bundles from their recent historical counterparts rises; and task-level prior-experience requirements increase without an analogous stable rise in schooling years. Outcome-specific samples, scales and inference procedures preclude ranking the four coefficients as if they measured the same investment.

More generally, the findings suggest a question for future research about the complementarity of different kinds of learning. A worker who gained a foundation in formal education may be better able to update an occupational practice when a new AI tool changes the feasible task set. Experience can then make a worker better at identifying errors, understanding context and coordinating with colleagues, while new tools may help the same worker apply those abilities to a broader range of activities. This interaction could increase demand for both a strong initial foundation and repeated, targeted updates even if an average advertised schooling threshold does not rise. Our data are consistent with some of these possible mechanisms, especially the changes in task composition and combination distance, but cannot rank them by their productivity effects. Future matched employer-worker data could compare entrants with different educational backgrounds within similar task environments and measure subsequent training, output quality and earnings. Such work would distinguish a durable change in the returns to investment from a temporary change in the wording of vacancies. Until then, the present results are best treated as a detailed description of changing demand signals alongside qualified evidence on their timing and exposure gradient.

The allocation of tasks and changes in their combinations are consistent with employers revising the relative value of existing knowledge and coordinated abilities. Formal education can still supply foundational, transferable skills and verifiable credentials. Practice, feedback and continuing learning may support adaptation when task weights and combinations change. These are implications for potential returns to investments; the vacancy data cannot establish that any worker has already changed her investment plan. Earlier research on science and technology finance and labour demand^[36] and on intelligent upgrading and labour costs^[37] addresses related resource-allocation mechanisms in other samples, rather than providing direct validation of our individual-level interpretation.

One practical response is to strengthen transferable reasoning, communication and verification in foundational and higher education. Short modular training linked to observable tasks could let workers update skills as assignments change; firms can use supervised projects, feedback on output quality and job rotation to develop task-integration experience. For entrants, internships and project-based learning can reduce the cost of acquiring a first credible record of work. Vacancy writers can distinguish experience genuinely required at

entry from abilities teachable after hiring. These proposals follow from the kinds of demand signals observed here, not from estimated treatment effects of any one educational policy.

Three limits remain central. Job descriptions do not reveal actual schooling purchases, training attendance, wages or realised returns. Outcome coverage changes across provinces, task mappings and job titles; the 1,783 identifiable combination cells represent 68.5% of the 2019–2025 potential cells. Several dynamic checks indicate anticipation or pre-existing trends, so December 2022 is a diffusion anchor rather than a universally exogenous break. Linking vacancies with worker training, mobility and wage records would allow future research to test the full pathway from task reorganisation through investment to labour-market returns.

There are also limits in the interpretation of an advertised experience requirement. A firm can use “years of experience” as a shorthand for reliability, knowledge of a workflow or reduced onboarding cost. An AI tool might substitute for some of these attributes and complement others. Thus a positive average requirement does not imply that every vacancy has become less accessible to a novice; the new-graduate and unrestricted-experience shares address adjacent but separate parts of that question. Further evidence would need to observe application and hiring rates for novices, the tasks they are assigned and whether assistance actually changes their productivity after hiring. Similar caution applies to schooling: an unchanged years threshold does not measure the content of a curriculum or the value of a specific field. These limitations are substantive reasons for the paper’s multidimensional outcome design, not minor caveats added after estimating the models.

An extension would join task-coded vacancies with establishment-level technology adoption and with subsequent hiring, tenure and wage records. This linkage could establish whether a changing bundle represents a new position or rewriting of a familiar advertisement, whether a listed requirement affects who is hired, and whether incumbent workers receive training before or after the change. Worker-level educational histories and training participation could then distinguish an incentive to update skills from a realised adjustment. A common-sample comparison of schooling and experience coefficients, together with an explicit account of missing combination cells, would strengthen the internal comparison of margins. These questions are left open

by the present design, which deliberately measures the employer side of a broader human-capital investment process.

Abbreviations

AI, artificial intelligence; GDP, gross domestic product; P2, second-stage task-combination measure.

References

- [1] Eloundou T, Manning S, Mishkin P, Rock D: GPTs are GPTs: An early look at the labor market impact potential of large language models. arXiv:2303.10130 (2023).
- [2] Noy S, Zhang W: Experimental evidence on the productivity effects of generative artificial intelligence. Science 381, 187–192 (2023). <https://doi.org/10.1126/science.adh2586>
- [3] Peng S, Kalliamvakou E, Cihon P, Demirer M: The impact of AI on developer productivity: Evidence from GitHub Copilot. arXiv:2302.06590 (2023).
- [4] Brynjolfsson E, Li D, Raymond L: Generative AI at work. Q. J. Econ. 140(2), 889–942 (2025). <https://doi.org/10.1093/qje/qjae044>
- [5] Dell’Acqua F, McFowland E, Mollick E R, et al.: Navigating the jagged technological frontier: Field experimental evidence of the effects of AI on knowledge worker productivity and quality. SSRN (2023). <https://doi.org/10.2139/ssrn.4573321>
- [6] Becker G S: Human capital: A theoretical and empirical analysis, with special reference to education, first edn. National Bureau of Economic Research, New York (1964).
- [7] Mincer J: Schooling, experience, and earnings. National Bureau of Economic Research, New York (1974).
- [8] Ben-Porath Y: The production of human capital and the life cycle of earnings. J. Polit. Econ. 75(4), 352–365 (1967).
- [9] Bick A, Blandin A, Deming D J: The rapid adoption of generative AI. NBER Working Paper 32966 (2024). <https://doi.org/10.3386/w32966>

824 [10] Lazear E P: Firm-specific human capital: A skill-weights approach. *J. Polit. Econ.* 117(5), 914–940
825 (2009). <https://doi.org/10.1086/648671>

826 [11] Acemoglu D, Autor D, Hazell J, et al.: Artificial intelligence and jobs: Evidence from online vacancies. *J.*
827 *Labor Econ.* 40(S1), S293–S340 (2022). <https://doi.org/10.1086/718327>

828 [12] Goldin C, Katz L F: *The race between education and technology*. Harvard University Press, Cambridge
829 (2008).

830 [13] Gibbons R, Waldman M: Task-specific human capital. *Am. Econ. Rev.* 94(2), 203–207 (2004).
831 <https://doi.org/10.1257/0002828041301579>

832 [14] Deming D, Kahn L B: Skill requirements across firms and labor markets: Evidence from job postings for
833 professionals. *J. Labor Econ.* 36(S1), S337–S369 (2018). <https://doi.org/10.1086/694106>

834 [15] Frey C B, Osborne M A: The future of employment: How susceptible are jobs to computerisation?
835 *Technol. Forecast. Soc. Change* 114, 254–280 (2017). <https://doi.org/10.1016/j.techfore.2016.08.019>

836 [16] Autor D H, Levy F, Murnane R J: The skill content of recent technological change: An empirical
837 exploration. *Q. J. Econ.* 118(4), 1279–1333 (2003).

838 [17] Autor D H: Why are there still so many jobs? The history and future of workplace automation. *J. Econ.*
839 *Perspect.* 29(3), 3–30 (2015). <https://doi.org/10.1257/jep.29.3.3>

840 [18] Acemoglu D, Restrepo P: Automation and new tasks: How technology displaces and reinstates labor. *J.*
841 *Econ. Perspect.* 33(2), 3–30 (2019). <https://doi.org/10.1257/jep.33.2.3>

842 [19] Acemoglu D, Restrepo P: The race between man and machine: Implications of technology for growth,
843 factor shares, and employment. *Am. Econ. Rev.* 108(6), 1488–1542 (2018).
844 <https://doi.org/10.1257/aer.20160696>

845 [20] Gmyrek P, Berg J, Kamiński K, et al.: Generative AI and jobs: A refined global index of occupational
846 exposure. ILO Working Paper 140 (2025).

847 [21] Agrawal A, Gans J S, Goldfarb A: Artificial intelligence: The ambiguous labor market impact of
848 automating prediction. *J. Econ. Perspect.* 33(2), 31–50 (2019). <https://doi.org/10.1257/jep.33.2.31>

849 [22] Chen L, Gao Y, Yu L: How artificial intelligence changes firms' labor demand: Evidence from large-
850 scale online recruitment data [in Chinese]. *Management World* (6), 74–93 (2024).

851 [23] He Q, Li X: Generative artificial intelligence and the differentiation of firm labor demand: A quasi-
852 natural experiment based on the release of ChatGPT [in Chinese]. *J. Shanghai Univ. Financ. Econ.* 28(4), 45–
853 59 (2026). <https://doi.org/10.16538/j.cnki.jsufe.2026.04.004>

854 [24] Humlum A, Vestergaard E: Still waters, rapid currents: Early labor market transformation under
855 generative AI. NBER Working Paper 33777 (2025). <https://doi.org/10.3386/w33777>

856 [25] Liao W, Guo Y, Zhao N: The impact of large language models on labor demand: Evidence from
857 recruitment data of listed companies [in Chinese]. *J. Financ. Econ.* 52(5), 109–123 (2026).
858 <https://doi.org/10.16538/j.cnki.jfe.20260331.401>

859 [26] Spitz-Oener A: Technical change, job tasks, and rising educational demands: Looking outside the wage
860 structure. *J. Labor Econ.* 24(2), 235–270 (2006). <https://doi.org/10.1086/499972>

861 [27] Deming D J: The growing importance of social skills in the labor market. *Q. J. Econ.* 132(4), 1593–1640
862 (2017). <https://doi.org/10.1093/qje/qjx022>

863 [28] Goos M, Manning A, Salomons A: Explaining job polarization: Routine-biased technological change and
864 offshoring. *Am. Econ. Rev.* 104(8), 2509–2526 (2014). <https://doi.org/10.1257/aer.104.8.2509>

865 [29] Atalay E, Phongthientham P, Sotelo S, Tannenbaum D: The evolution of work in the United States. *Am.*
866 *Econ. J. Appl. Econ.* 12(2), 1–34 (2020). <https://doi.org/10.1257/app.20190070>

867 [30] Hyman B G, Lahey B, Ni K, Pilossoph L: How retrainable are AI-exposed workers? NBER Working
868 Paper 34174 (2025). <https://doi.org/10.3386/w34174>

869 [31] Card D, Kluve J, Weber A: What works? A meta analysis of recent active labor market program
870 evaluations. *J. Eur. Econ. Assoc.* 16(3), 894–931 (2018). <https://doi.org/10.1093/jeea/jvx028>

871 [32] Autor D H, Handel M J: Putting tasks to the test: Human capital, job tasks, and wages. *J. Labor Econ.*
872 31(S1), S59–S96 (2013). <https://doi.org/10.1086/669332>

- 873 [33] Felten E, Raj M, Seamans R: Occupational, industry, and geographic exposure to artificial intelligence: A
874 novel dataset and its potential uses. *Strateg. Manag. J.* 42(12), 2195–2217 (2021).
875 <https://doi.org/10.1002/smj.3286>
- 876 [34] Gmyrek P, Berg J, Bescond D: Generative AI and jobs: A global analysis of potential effects on job
877 quantity and quality. ILO Working Paper 96 (2023).
- 878 [35] Aitchison J: The statistical analysis of compositional data. *J. R. Stat. Soc. Series B* 44(2), 139–160 (1982).
879 <https://doi.org/10.1111/j.2517-6161.1982.tb01195.x>
- 880 [36] Zhou J, Guo K: The development of science and technology finance and the labor demand of high-tech
881 enterprises [in Chinese]. *Technical Economics* 44(3), 15–28 (2025). [https://doi.org/10.12404/j.issn.1002-](https://doi.org/10.12404/j.issn.1002-980X.J24121102)
882 [980X.J24121102](https://doi.org/10.12404/j.issn.1002-980X.J24121102)
- 883 [37] Liang Y, Huang X: The dynamic impact of enterprise intelligence upgrade on labor costs [in Chinese].
884 *Technical Economics* 44(2), 115–129 (2025). <https://doi.org/10.12404/j.issn.1002-980X.J24032302>

885 **Declarations**

886 **Availability of data and materials**

887 The underlying job-posting microdata are subject to the original data provider’s access conditions and are not
888 redistributed with this article. The definitions of the constructed measures and the aggregate results are
889 reported in the article and supplementary tables. Requests concerning access to the underlying data or
890 reproducibility materials may be directed to the corresponding author; any sharing is subject to the applicable
891 data-use terms.

892 **Competing interests**

893 The author declares that there are no competing interests.

894 **Funding**

895 This research received no external funding and was not supported by a funded research project.

896 **Authors’ contributions**

897 Yuan Dong is the sole author and was responsible for the study design, data preparation and analysis,
898 interpretation of the results, and drafting and revising the manuscript. The author approved the final version.

899 Acknowledgements

900 Not applicable.

901 Appendix

902 The appendix reports stage-specific effects, alternative shocks, identification diagnostics, regional
 903 heterogeneity, seven pre-entry human-capital measures, and valid-posting threshold checks. Each table
 904 specifies its estimation unit, standard errors, and fixed effects.

905 **Table A1. Stage-specific task reallocation effects**

Variable	Absolute reallocation contribution($\times 10^{-6}$)	Expansion contribution($\times 10^{-6}$)	Contraction contribution($\times 10^{-6}$)
Early $\times$ exposure quartile 2	-0.608 (0.662)	-0.598 (0.805)	-0.618 (1.720)
Early $\times$ exposure quartile 3	-0.210 (0.569)	-0.163 (0.584)	-0.257 (1.180)
Early $\times$ exposure quartile 4	0.346 (0.672)	-1.000 (0.712)	1.690 (1.270)
Middle $\times$ exposure quartile 2	0.067 (0.741)	1.130 (1.380)	-0.994 (1.600)
Middle $\times$ exposure quartile 3	-0.258 (0.717)	-1.430 (1.310)	0.909 (1.140)
Middle $\times$ exposure quartile 4	-0.011 (0.725)	-1.920 (1.510)	1.900 (1.750)
Late $\times$ exposure quartile 2	-1.430 (0.876)	-0.402 (1.110)	-2.450* (1.300)
Late $\times$ exposure quartile 3	-1.480* (0.786)	-1.040 (0.674)	-1.910 (1.430)
Late $\times$ exposure quartile 4	-0.833 (0.779)	-1.530 (1.360)	-0.133 (1.400)
Constant	83.700*** (0.000)	83.700*** (0.000)	83.700*** (0.000)
Observations	2364205	2364205	2364205
R^2	0.903	0.896	0.877
Adjusted R^2	0.897	0.890	0.869
Controls	Absorbed by province $\times$ month fixed effects	Absorbed by province $\times$ month fixed effects	Absorbed by province $\times$ month fixed effects
Province fixed effects	Province $\times$ task	Province $\times$ task	Province $\times$ task
Time fixed effects	Province $\times$ month; task $\times$ month	Province $\times$ month; task $\times$ month	Province $\times$ month; task $\times$ month
Province-specific trends	No	No	No
Other high-dimensional fixed effects	Yes	Yes	Yes

906 *Note: Standard errors clustered by province and task are in parentheses. Coefficients and standard*
 907 *errors are multiplied by 10^6 ; observations and R^2 are unscaled. *, **, and *** denote significance*
 908 *at 10%, 5%, and 1%, respectively, using the original estimates. Stage definitions follow the source*
 909 *output and are not a monthly event study.*

910 **Table A2. Alternative shock measures: conventional demand share**

Variable	Conventional demand share($\times 10^{-4}$)	Conventional demand share($\times 10^{-4}$)	Conventional demand share($\times 10^{-4}$)
Regional GAI exposure shock	0.371 (2.680)		
Regional GAI exposure shock (log alternative)		-0.586 (2.420)	
Regional GAI exposure shock (post-shock alternative)			-2.810 (4.420)
Constant	400.000*** (0.271)	400.000*** (0.249)	400.000*** (0.206)
Observations	2184097	2184097	2184097

Variable	Conventional demand share($\times 10^{-4}$)	Conventional demand share($\times 10^{-4}$)	Conventional demand share($\times 10^{-4}$)
R^2	0.431	0.431	0.431
Adjusted R^2	0.275	0.275	0.275
Controls	No	No	No
Province fixed effects	Province $\times$ standardized title	Province $\times$ standardized title	Province $\times$ standardized title
Time fixed effects	Calendar month	Calendar month	Calendar month
Province-specific trends	No	No	No
Other high-dimensional fixed effects	Yes	Yes	Yes

*Note: Standard errors clustered by province are in parentheses. Coefficients and standard errors are multiplied by 10^4 ; observations and R^2 are unscaled. *, **, and *** denote significance at 10%, 5%, and 1%, respectively.*

Table A3. Alternative shock measures: conventional demand intensity

Variable	Conventional demand intensity($\times 10^{-4}$)	Conventional demand intensity($\times 10^{-4}$)	Conventional demand intensity($\times 10^{-4}$)
Regional GAI exposure shock	1.340 (3.000)		
Regional GAI exposure shock (log alternative)		0.203 (2.620)	
Regional GAI exposure shock (post- shock alternative)			-1.570 (4.750)
Constant	430.000*** (0.303)	430.000*** (0.268)	430.000*** (0.221)
Observations	2184097	2184097	2184097
R^2	0.433	0.433	0.433
Adjusted R^2	0.277	0.277	0.277
Controls	No	No	No
Province fixed effects	Province $\times$ standardized title	Province $\times$ standardized title	Province $\times$ standardized title
Time fixed effects	Calendar month	Calendar month	Calendar month
Province-specific trends	No	No	No
Other high-dimensional fixed effects	Yes	Yes	Yes

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A4. Province-specific linear trends: task-combination distance

Variable	Task-combination distance	Task-combination distance
Regional GAI exposure shock	0.079*** (0.028)	0.086* (0.044)
Constant	-0.138*** (0.003)	-10.661*** (0.715)
Observations	1783	1783
R^2	0.215	0.251
Adjusted R^2	0.162	0.185
Controls	No	No
Province fixed effects	Yes	Yes
Time fixed effects	Yes	Yes
Province-specific trends	No	Yes
Other high-dimensional fixed effects	No	No

*Note: Standard errors are in parentheses; *, **, and *** indicate significance at the 10%, 5%, and 1% levels. Coefficients and standard errors are displayed to three decimal places. Province-month standard errors are clustered by province. Column (2) adds province-specific linear time trends; their individual coefficients and standard errors are omitted for brevity.*

Table A5. Alternative shock measures: task-combination distance

Variable	Task-combination distance	Task-combination distance	Task-combination distance
Regional GAI exposure shock	0.079*** (0.028)		

Variable	Task-combination distance	Task-combination distance	Task-combination distance
Regional GAI exposure shock (log alternative)		0.077** (0.030)	
Regional GAI exposure shock (post-shock alternative)			0.145** (0.067)
Constant	-0.138*** (0.003)	-0.137*** (0.004)	-0.137*** (0.004)
Observations	1783	1783	1783
R^2	0.215	0.215	0.214
Adjusted R^2	0.162	0.161	0.160
Controls	No	No	No
Province fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Province-specific trends	No	No	No
Other high-dimensional fixed effects	No	No	No

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A6. Identification checks for conventional demand

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
6-month lead of regional GAI exposure shock	1.410×10^{-4} (0.001)		
Log real GDP per capita	0.008 (0.011)	0.009 (0.011)	0.007 (0.010)
Fiscal self-sufficiency	-8.000×10^{-6} (0.002)	3.240×10^{-4} (0.002)	-3.230×10^{-4} (0.002)
Trade openness	-2.860×10^{-4} *** (1.000×10^{-4})	-2.670×10^{-4} *** (9.800×10^{-5})	-2.720×10^{-4} *** (9.400×10^{-5})
Child dependency ratio	-2.760×10^{-4} (4.480×10^{-4})	-2.290×10^{-4} (4.550×10^{-4})	-2.210×10^{-4} (4.040×10^{-4})
Old-age dependency ratio	-3.970×10^{-4} (0.001)	-3.690×10^{-4} (0.001)	-4.560×10^{-4} (0.001)
Tertiary-sector share	0.001 (0.019)	0.001 (0.019)	0.002 (0.016)
12-month lead of regional GAI exposure shock		1.320×10^{-4} (0.001)	
Capability-frontier Bartik alternative shock			-0.025** (0.010)
Constant	-0.010 (0.091)	-0.020 (0.092)	-0.004 (0.083)
Observations	2484	2424	2542
R^2	0.221	0.224	0.217
Controls	Yes	Yes	Yes
Province fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Province-specific trends	No	No	No
Other high-dimensional fixed effects	No	No	No

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A7. Identification checks for task reallocation

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
6-month lead of regional GAI exposure shock	0.005*** (0.001)		
Log real GDP per capita	0.001	-0.003	-2.330×10^{-4}

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
	(0.015)	(0.015)	(0.020)
Fiscal self-sufficiency	-0.002 (0.002)	-0.002 (0.002)	-0.002 (0.002)
Trade openness	6.800×10^{-5} (3.320×10^{-4})	3.400×10^{-5} (3.140×10^{-4})	1.670×10^{-4} (3.610×10^{-4})
Child dependency ratio	0.004*** (0.001)	0.004*** (0.001)	0.004*** (0.001)
Old-age dependency ratio	-0.001* (0.001)	-0.001 (0.001)	-0.002* (0.001)
Tertiary-sector share	0.031 (0.094)	0.047 (0.099)	-0.093 (0.086)
12-month lead of regional GAI exposure shock		0.005*** (0.001)	
Capability-frontier Bartik alternative shock			0.094 (0.070)
Constant	0.085 (0.167)	0.107 (0.165)	0.106 (0.214)
Observations	2100	2040	2158
R ²	0.891	0.891	0.889
Controls	Yes	Yes	Yes
Province fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Province-specific trends	No	No	No
Other high-dimensional fixed effects	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A8. Identification checks for task combinations

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
6-month lead of regional GAI exposure shock	0.002 (0.001)		
Log real GDP per capita	-0.009 (0.023)	-0.013 (0.028)	-0.005 (0.024)
Fiscal self-sufficiency	0.007** (0.003)	0.007** (0.003)	0.007* (0.003)
Trade openness	-2.800×10^{-5} (1.730×10^{-4})	1.800×10^{-5} (1.790×10^{-4})	3.900×10^{-5} (1.670×10^{-4})
Child dependency ratio	-0.001 (0.001)	-0.001 (0.001)	-4.020×10^{-4} (0.001)
Old-age dependency ratio	2.600×10^{-4} (0.002)	3.890×10^{-4} (0.002)	2.860×10^{-4} (0.002)
Tertiary-sector share	-0.100 (0.061)	-0.103 (0.062)	-0.110* (0.059)
12-month lead of regional GAI exposure shock		0.001 (0.001)	
Capability-frontier Bartik alternative shock			0.008 (0.023)
Constant	0.677*** (0.227)	0.710** (0.265)	0.636** (0.236)
Observations	2447	2387	2505
R ²	0.235	0.235	0.233
Controls	Yes	Yes	Yes
Province fixed effects	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes
Province-specific trends	No	No	No
Other high-dimensional fixed effects	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero

values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A9. Identification checks for education requirements

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
Triple-difference lead: 6 months	0.001 (0.002)		
Triple-difference lead: 12 months		0.002 (0.002)	
Triple-difference effect: Bartik alternative shock			0.080** (0.034)
Constant	15.403*** (4.100×10 ⁻⁵)	15.401*** (2.500×10 ⁻⁵)	15.418*** (0.002)
Observations	4217299	3830169	4623068
R ²	0.740	0.746	0.735
Controls	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects
Province fixed effects	Province × task	Province × task	Province × task
Time fixed effects	Province × month; task × month	Province × month; task × month	Province × month; task × month
Province-specific trends	No	No	No
Other high-dimensional fixed effects	Yes	Yes	Yes

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A10. Identification checks for experience requirements

Variable	6-month lead	12-month lead	Capability-frontier Bartik alternative shock
Triple-difference lead: 6 months	0.003*** (0.001)		
Triple-difference lead: 12 months		0.004*** (0.001)	
Triple-difference effect: Bartik alternative shock			0.033 (0.047)
Constant	2.686*** (2.300×10 ⁻⁵)	2.668*** (1.500×10 ⁻⁵)	2.696*** (0.004)
Observations	4055270	3680797	4451852
R ²	0.630	0.636	0.622
Controls	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects
Province fixed effects	Province × task	Province × task	Province × task
Time fixed effects	Province × month; task × month	Province × month; task × month	Province × month; task × month
Province-specific trends	No	No	No
Other high-dimensional fixed effects	Yes	Yes	Yes

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A11. Regional heterogeneity in task reallocation

Variable	Task-reallocation demand
Eastern × regional GAI exposure shock	0.011*** (0.002)
Central × regional GAI exposure shock	0.004***

Variable	Task-reallocation demand
	(0.001)
Western × regional GAI exposure shock	0.002 (0.001)
Northeastern × regional GAI exposure shock	0.020 (0.013)
Log real GDP per capita	0.003 (0.019)
Fiscal self-sufficiency	-0.001 (0.002)
Trade openness	7.300×10^{-5} (3.500×10^{-4})
Child dependency ratio	0.003*** (0.001)
Old-age dependency ratio	-0.001 (0.001)
Tertiary-sector share	0.013 (0.104)
Constant	0.082 (0.203)
Observations	2158
R^2	0.891
Controls	Yes
Province fixed effects	Yes
Time fixed effects	Yes
Province-specific trends	No
Other high-dimensional fixed effects	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A12. Regional heterogeneity in task combinations

Variable	Task-combination demand
Eastern × regional GAI exposure shock	0.002 (0.002)
Central × regional GAI exposure shock	-0.003 (0.003)
Western × regional GAI exposure shock	0.008*** (0.002)
Northeastern × regional GAI exposure shock	-0.040* (0.020)
Log real GDP per capita	-0.004 (0.022)
Fiscal self-sufficiency	0.007* (0.003)
Trade openness	-1.600×10^{-5} (1.540×10^{-4})
Child dependency ratio	7.400×10^{-5} (0.001)
Old-age dependency ratio	-9.000×10^{-6} (0.002)
Tertiary-sector share	-0.101* (0.059)
Constant	0.618*** (0.216)
Observations	2505
R^2	0.240
Controls	Yes
Province fixed effects	Yes
Time fixed effects	Yes
Province-specific trends	No
Other high-dimensional fixed effects	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A13. Regional heterogeneity in conventional demand

Variable	Conventional-task demand
Eastern × regional GAI exposure shock	0.001 (0.001)
Central × regional GAI exposure shock	-0.001 (0.001)
Western × regional GAI exposure shock	0.001 (0.003)
Northeastern × regional GAI exposure shock	-0.014* (0.007)
Log real GDP per capita	0.008 (0.009)
Fiscal self-sufficiency	-2.850×10 ⁻⁴ (0.002)
Trade openness	-2.780×10 ⁻⁴ *** (9.700×10 ⁻⁵)
Child dependency ratio	-2.280×10 ⁻⁴ (4.050×10 ⁻⁴)
Old-age dependency ratio	-0.001 (0.001)
Tertiary-sector share	0.003 (0.018)
Constant	-0.014 (0.078)
Observations	2542
R ²	0.217
Controls	Yes
Province fixed effects	Yes
Time fixed effects	Yes
Province-specific trends	No
Other high-dimensional fixed effects	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A14. Structural heterogeneity in education requirements

Variable	Tertiary sector	Fiscal self-sufficiency	Trade openness
Continuous triple interaction	0.001 (0.003)	-0.001 (0.004)	0.001 (0.003)
Main effect × tertiary-sector share	0.005*** (0.001)		
Main effect × fiscal self-sufficiency		0.008** (0.003)	
Main effect × trade openness			0.006** (0.003)
Constant	15.412*** (3.800×10 ⁻⁵)	15.412*** (1.090×10 ⁻⁴)	15.412*** (1.320×10 ⁻⁴)
Observations	4623068	4623068	4623068
R ²	0.735	0.735	0.735
Controls	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects	Absorbed by province × month fixed effects
Province fixed effects	Province × task	Province × task	Province × task
Time fixed effects	Province × month; task × month	Province × month; task × month	Province × month; task × month
Province-specific trends	No	No	No

Variable	Tertiary sector	Fiscal self-sufficiency	Trade openness
Other high-dimensional fixed effects	Yes	Yes	Yes

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A15. Regional heterogeneity in education requirements

Variable	Mean years of education
Eastern × triple-difference effect	0.011 (0.007)
Central × triple-difference effect	-4.300×10 ⁻⁴ (0.003)
Western × triple-difference effect	-0.002 (0.007)
Northeastern × triple-difference effect	-0.014 (0.030)
Constant	15.412*** (1.550×10 ⁻⁴)
Observations	4623068
R ²	0.735
Controls	Absorbed by province × month fixed effects
Province fixed effects	Province × task
Time fixed effects	Province × month; task × month
Province-specific trends	No
Other high-dimensional fixed effects	Yes

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A16. Regional heterogeneity in experience requirements

Variable	Mean years of experience
Eastern × triple-difference effect	0.008* (0.004)
Central × triple-difference effect	0.007** (0.003)
Western × triple-difference effect	0.011*** (0.004)
Northeastern × triple-difference effect	0.037 (0.023)
Constant	2.693*** (1.030×10 ⁻⁴)
Observations	4451852
R ²	0.622
Controls	Absorbed by province × month fixed effects
Province fixed effects	Province × task
Time fixed effects	Province × month; task × month
Province-specific trends	No
Other high-dimensional fixed effects	Yes

*Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.*

Table A17. Control specifications: mean years of education

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	-0.015** (0.007)	-0.014* (0.007)	-0.015 (0.010)	-0.015 (0.010)
Log real GDP per capita		0.021 (0.113)	-0.038 (0.107)	-0.047 (0.107)
Fiscal self-sufficiency		0.027 (0.016)	0.022 (0.016)	0.023 (0.016)
Trade openness		-0.001 (0.002)	-0.001 (0.002)	-0.001 (0.002)
Child dependency ratio			-0.012 (0.007)	-0.011 (0.008)
Old-age dependency ratio			0.004 (0.010)	-0.003 (0.011)
Tertiary-sector share			-0.270 (0.378)	-0.302 (0.380)
Urbanization rate				-0.639 (0.480)
Constant	15.330*** (0.000)	15.152*** (1.007)	16.017*** (0.985)	16.632*** (1.070)
Observations	2964	2914	2542	2542
R^2	0.601	0.641	0.642	0.642
Adjusted R^2	0.584	0.625	0.624	0.625
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A18. Control specifications: bachelor's degree or above

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	-0.001 (0.003)	-0.001 (0.003)	0.005 (0.004)	0.004 (0.004)
Log real GDP per capita		0.016 (0.039)	-0.018 (0.047)	-0.024 (0.046)
Fiscal self-sufficiency		0.015** (0.007)	0.013* (0.007)	0.014* (0.007)
Trade openness		-3.370×10 ⁻⁴ (0.001)	-0.001 (0.001)	-0.001 (0.001)
Child dependency ratio			-0.003 (0.003)	-0.003 (0.003)
Old-age dependency ratio			0.001 (0.004)	-0.003 (0.004)
Tertiary-sector share			-0.146 (0.124)	-0.167 (0.115)
Urbanization rate				-0.408* (0.213)
Constant	0.536*** (0.000)	0.397 (0.343)	0.847* (0.442)	1.240*** (0.420)
Observations	2964	2914	2542	2542
R^2	0.564	0.611	0.618	0.619
Adjusted R^2	0.545	0.594	0.599	0.601
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A19. Control specifications: graduate degree or above

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	3.150×10^{-4} (0.001)	0.001 (0.001)	-4.360×10^{-4} (0.001)	-0.001 (0.001)
Log real GDP per capita		-0.017 (0.015)	-0.007 (0.014)	-0.008 (0.013)
Fiscal self-sufficiency		0.001 (0.001)	0.001 (0.001)	0.002 (0.001)
Trade openness		-3.300×10^{-4} *** (1.600×10^{-4})	-2.980×10^{-5} (1.400×10^{-4})	-2.480×10^{-5} (1.440×10^{-4})
Child dependency ratio			-3.170×10^{-4} (0.001)	-2.680×10^{-4} (0.001)
Old-age dependency ratio			0.001 (0.001)	0.001 (0.001)
Tertiary-sector share			-0.012 (0.033)	-0.017 (0.029)
Urbanization rate				-0.085** (0.041)
Constant	0.030*** (0.000)	0.185 (0.130)	0.074 (0.121)	0.156 (0.109)
Observations	2964	2914	2542	2542
R ²	0.286	0.479	0.441	0.443
Adjusted R ²	0.254	0.455	0.414	0.416
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A20. Control specifications: unrestricted education

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	0.002 (0.002)	0.003* (0.002)	0.001 (0.004)	0.001 (0.004)
Log real GDP per capita		-0.031 (0.027)	-0.035 (0.030)	-0.035 (0.030)
Fiscal self-sufficiency		-0.003 (0.004)	-0.005 (0.004)	-0.005 (0.004)
Trade openness		7.380×10^{-6} (2.010×10^{-4})	2.010×10^{-4} (3.530×10^{-4})	2.010×10^{-4} (3.520×10^{-4})
Child dependency ratio			0.004 (0.003)	0.004 (0.003)
Old-age dependency ratio			-0.004 (0.003)	-0.004 (0.004)
Tertiary-sector share			-0.043 (0.124)	-0.043 (0.122)
Urbanization rate				0.002 (0.189)
Constant	0.095*** (0.000)	0.370 (0.235)	0.397 (0.268)	0.395 (0.357)
Observations	2964	2914	2542	2542
R ²	0.638	0.639	0.622	0.622

Variable	Baseline	Core controls	Basic controls	Extended controls
Adjusted R^2	0.622	0.622	0.604	0.604
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A21. Control specifications: mean years of experience

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	0.006 (0.013)	0.005 (0.012)	0.012 (0.018)	0.011 (0.018)
Log real GDP per capita		-0.043 (0.183)	0.027 (0.174)	0.021 (0.172)
Fiscal self-sufficiency		0.042 (0.032)	0.059* (0.034)	0.059* (0.034)
Trade openness		-0.003 (0.002)	-0.003 (0.002)	-0.003 (0.002)
Child dependency ratio			-0.007 (0.009)	-0.007 (0.009)
Old-age dependency ratio			0.016* (0.009)	0.012 (0.012)
Tertiary-sector share			0.849** (0.339)	0.829** (0.343)
Urbanization rate				-0.385 (0.625)
Constant	2.438*** (0.000)	2.879* (1.585)	1.645 (1.504)	2.015 (1.427)
Observations	2962	2914	2542	2542
R^2	0.454	0.496	0.515	0.515
Adjusted R^2	0.429	0.473	0.491	0.491
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A22. Control specifications: new-graduate eligibility

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	0.001 (0.002)	0.002 (0.002)	0.002 (0.004)	0.002 (0.004)
Log real GDP per capita		-0.025 (0.021)	-0.025 (0.025)	-0.026 (0.025)
Fiscal self-sufficiency		-0.007 (0.005)	-0.010* (0.005)	-0.010* (0.005)
Trade openness		4.300×10^{-4} * (2.420×10^{-4})	0.001 ** (2.840×10^{-4})	0.001 ** (2.840×10^{-4})
Child dependency ratio			0.002 (0.001)	0.002 (0.001)
Old-age dependency ratio			-0.003 (0.002)	-0.003 (0.002)
Tertiary-sector share			0.033	0.030

Variable	Baseline	Core controls	Basic controls	Extended controls
			(0.048)	(0.050)
Urbanization rate				-0.057
				(0.082)
Constant	0.065***	0.275	0.248	0.303
	(0.000)	(0.186)	(0.236)	(0.252)
Observations	2964	2914	2542	2542
R ²	0.543	0.627	0.625	0.626
Adjusted R ²	0.522	0.610	0.607	0.607
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A23. Control specifications: unrestricted experience

Variable	Baseline	Core controls	Basic controls	Extended controls
Regional GAI exposure shock	0.002	0.002	0.001	0.001
	(0.003)	(0.003)	(0.005)	(0.005)
Log real GDP per capita		-0.006	-0.002	0.003
		(0.032)	(0.028)	(0.027)
Fiscal self-sufficiency		-0.002	-0.003	-0.003
		(0.005)	(0.006)	(0.005)
Trade openness		-3.690×10 ⁻⁴	-2.710×10 ⁻⁵	-4.640×10 ⁻⁵
		(2.980×10 ⁻⁴)	(4.630×10 ⁻⁴)	(4.460×10 ⁻⁴)
Child dependency ratio			0.003	0.003
			(0.002)	(0.002)
Old-age dependency ratio			0.001	0.004
			(0.003)	(0.004)
Tertiary-sector share			-0.107	-0.091
			(0.126)	(0.115)
Urbanization rate				0.326**
				(0.135)
Constant	0.213***	0.272	0.207	-0.106
	(0.000)	(0.276)	(0.252)	(0.257)
Observations	2964	2914	2542	2542
R ²	0.691	0.728	0.702	0.703
Adjusted R ²	0.677	0.716	0.687	0.688
Controls	No	Core controls	Basic controls	Extended controls
Province fixed effects	Yes	Yes	Yes	Yes
Time fixed effects	Yes	Yes	Yes	Yes
Province-specific trends	No	No	No	No
Other high-dimensional fixed effects	No	No	No	No

Note: Standard errors are in parentheses. *, **, and *** indicate significance at the 10%, 5%, and 1% levels, respectively, based on unrounded estimates. Unless the heading specifies a scale factor, coefficients and standard errors are shown to at most three decimal places; very small nonzero values use scientific notation. Observations are integers with thin-space digit grouping. Task-level standard errors are clustered by province and task; province-month regressions are clustered by province.

Table A24. Minimum valid-posting thresholds for task combinations

Measure and threshold	Observations	Coefficient (standard error)	Conventional p-value	Wild-bootstrap p-value	Wild-bootstrap 95% CI
Distance: original sample	1783	0.0787 (0.0277)	0.0080	0.0252	[0.0177, 0.1528]
Distance: ≥5 valid postings	1488	0.0775 (0.0218)	0.0014	0.0193	[0.0210, 0.1362]
Distance: ≥10 valid postings	1293	0.0746 (0.0163)	0.0001	0.0020	[0.0459, 0.1377]
Composite: original sample	1783	0.0470 (0.0219)	0.0402	0.1146	[-0.0268, 0.1005]

Composite: ≥ 5 valid postings	1488	0.0501 (0.0157)	0.0036	0.0704	[-0.0074, 0.0777]
Composite: ≥ 10 valid postings	1293	0.0536 (0.0141)	0.0008	0.0113	[0.0155, 0.0811]

Note: “Distance” denotes task-combination distance and “Composite” the composite task-combination measure. Coefficients use the province and calendar-month fixed-effects specification of Table 6, with province-clustered standard errors in parentheses. Wild-bootstrap p-values and 95% confidence intervals follow the province-clustered FWL procedure of Table 7 (9,999 replications; seed 20260919). The thresholds restrict the number of valid postings used to construct the measures; the original sample has no additional threshold. Observations are the effective regression sample. The original, five-posting, and ten-posting thresholds cover 31, 29, and 28 provinces, respectively. Coefficients, standard errors, and p-values are shown to four decimal places; significance uses unrounded results. This is a threshold and inference summary; its model fit measures are not provided in the frozen threshold results and have not been imputed.