The Broken Ladder: AI, Remote Work, and Early-Career Hiring

Peter John Lambert[†] & Yannick Schindler[‡]

May 2026[*]

Abstract

Is generative AI replacing junior workers? A growing literature answers yes, citing large declines in early-career hiring concentrated in GenAI-exposed occupations. We argue that this verdict is premature because GenAI exposure is strongly correlated with another post-pandemic shock, working from home (WFH). Using two data sources spanning 243 million new hires and 407 million online job postings, collected across the US, UK, Canada, and Australia during 2017-2025, we estimate difference-in-difference designs at the occupation, region, and firm level. When estimated separately, a two-standard-deviation increase in GenAI and WFH exposure each predicts, by 2025, a fall of around 5pp in the junior-share of new hires and around 3pp in the share of job ads requiring limited experience. Estimated jointly, the WFH effect remains, while the GenAI coefficient attenuates sharply and is often statistically indistinguishable from zero. Alternative exposure measures, residualization designs, flexible non-parametric co-treatment controls, and replacing exposure-measures with actual WFH adoption as the treatment all support our finding that WFH is a robust predictor of the decline in early-career hiring.

Keywords: Remote Work, Generative AI, Junior Hiring, Technology Exposure, Omitted Variable Bias
JEL Codes: J23, J24, J31, O33

[*†] University of Warwick & London School of Economics, [‡] Ellison Institute of Technology, Oxford. Correspondence: p.j.lambert@lse.ac.uk. We thank John Van Reenen, Tim Besley, Mirko Draca, Arnaud Dyèvre, Steve Pischke, Nick Bloom, Andrew Scott, Steven Davis, Bledi Taska, Raffaella Sadun, and Stephen Hansen. We also wish to thank seminar participants at LSE, University of Warwick, Ellison Institute of Technology (EIT), Cambridge University, London Business School, the ESRC Executive Committee Workshop, and WZB Berlin. We acknowledge financial support from the ESRC via CAGE (University of Warwick) and the Centre for Economic Performance (CEP) at LSE. Views expressed are those of the authors and not those of the institutions with which they are affiliated.

1 Introduction

Across many developed countries, the share of new hires going to early-career workers has fallen sharply since 2022 (Figure 1). The implication is stark. Early-career jobs are the principal channel through which workers accumulate firm-specific and general human capital (Arrow, 1962; Mincer, 1974). A persistent contraction of this kind hollows out the pipeline of future experienced workers, causing declines in aggregate productivity as well as imposing cohort-specific scarring (Pallais, 2014; Oreopoulos et al., 2012; Schwandt and von Wachter, 2019). A new literature lays the blame on generative AI tools, arguing that GenAI now substitutes directly for the cognitive and routine-analytical tasks firms used to assign to early-career workers (Brynjolfsson et al., 2025a; Teeselink, 2025; Azar et al., 2025; Hosseini Maasoum and Lichtinger, 2025).

This paper argues that this verdict is premature. Many of the empirical designs employed to support the claim that GenAI is driving declines in the attractiveness of hiring early-career workers face a first-order confounding factor: the rapid and persistent shift to WFH arrangements (Barrero et al., 2021; Hansen et al., 2023). Two facts support the hypothesis that WFH is an important omitted variable:

First, each technology has a plausible mechanism for disrupting junior demand. GenAI may automate a substantial fraction of the cognitive and routine-analytical tasks that have typically been delegated to juniors (Noy and Zhang, 2023; Eloundou et al., 2024; Brynjolfsson et al., 2025a), though how mechanically static task-level exposure maps to hiring is contested (Demirer et al., 2026; Gans and Goldfarb, 2026; Garicano et al., 2026). WFH has been shown to raise the cost of supervising and monitoring workers, and can slow on-the-job learning (Emanuel et al., 2024; Emanuel et al., 2026; Yang et al., 2022; Aksoy et al., 2026). These organizational frictions can erode the value-proposition of investing in early-career talent.

Second, the exposure to GenAI and WFH fall on largely the same occupations: white-collar, knowledge-intensive roles with heavy computer use. At the 6-digit O*NET-SOC occupation level, the WFH exposure measure of Hansen et al. (2023) and the GenAI exposure index of Eloundou et al. (2024) have a Spearman rank correlation of 0.77 (Figure 2). Other measures of both GenAI and WFH exposure are similarly correlated. Software developers, technical writers, and management consultants sit at the top of both rankings, while electricians, janitors, and construction laborers sit at the bottom.

The empirical analysis in this paper exploits the timing and differential incidence of both shocks using two separate microdata sets covering the United States, the United Kingdom, Canada, and Australia. Our first dataset is constructed from 243 million employer–employee linked records

of new hires assembled from resume data, which we use to measure the share of all new hires going to junior workers.[1] The second panel comprises 407 million online job vacancy postings collected across the same countries, and is used to measure the share of postings requiring no more than three years of relevant work experience. We combine this with widely used measures of GenAI and WFH exposure.

Our difference-in-differences estimation strategy exploits the timing of WFH and GenAI adoption, together with the continuous exposure measures above, which vary across occupations. We use this to identify the effects of each technology on early-career hiring and recruitment, both at the firm level and at the occupation-region level. In single-treatment difference-in-differences specifications, both WFH and GenAI exposure predict sizable declines in early-career job opportunities in 2023–2025. A two-standard-deviation increase in WFH or GenAI exposure is associated with a fall of around 4–5 pp in the junior share of new hires by 2025, and around 3 pp in the share of postings requiring limited experience.

Taken individually, both shocks look like plausible drivers of the junior-hiring decline. Our first main result comes from a joint-treatment specification in which both exposures enter simultaneously. The WFH effects are robust, and depict a similar event-study coefficient path to the single-treatment results. By contrast, the GenAI effects attenuate heavily and often become statistically indistinguishable from zero. This pattern is consistent with the hypothesis that WFH is a key omitted variable in the single-treatment GenAI specifications, and that the junior-hiring decline is more strongly driven by WFH than by GenAI.

Several robustness exercises probe whether this asymmetry is an artifact of functional form, multi-collinearity, or measurement error. A semi-parametric exercise in the spirit of Robinson (1988) absorbs the co-treatment exposure via flexible step functions of progressively finer quantile resolution. This reveals that the WFH coefficient is stable to highly flexible forms of GenAI exposure, but that a single WFH dummy variable is enough to render the GenAI coefficient statistically insignificant. Using alternative WFH and GenAI exposure measures, residualized single-treatment designs also support the fact that WFH is a strong predictor of our outcomes of interest, whereas GenAI exposure is not. We also conduct standard multi-collinearity and selection-on-unobservables diagnostics (Altonji et al., 2008; Oster, 2019; Cinelli and Hazlett, 2020), and calibrated Monte Carlo simulations of classical measurement error all preserve the WFH-dominant ranking. The Monte Carlo exercise in particular shows that implausibly high

[1] 'Junior' is measured using the position being filled and the worker-history on their resumes, see Section 4 for details.

mismeasurement of GenAI exposure would be required to spuriously misattribute a true GenAI effect to WFH.

A natural question is whether the exposure-based empirical designs are reliable signals of the impact of real adoption of WFH arrangements. Our final substantive contribution is to test this question head on. We use a design where treatment assignment is based on *actual* WFH adoption at the firm and region-level. In this test, we first match firms or regions based on the propensity to adopt WFH/GenAI implied by their occupation-mix, and then assess the effect of posting job ads which explicitly offered remote/hybrid work arrangements in 2021–2022 on the relative hiring and recruitment of early-career workers in 2023-2025. We find that actual WFH adoption predicts a fall in the junior-share of new hires and the share of job postings requiring 3 years of experience or less. The magnitudes of these effects are similar to our other findings, consistent with our hypothesis that WFH has tilted hiring and recruitment towards more senior and more experienced workers.

To put our main findings on the effect of WFH on early-career hiring in concrete terms, consider two illustrative occupations: construction carpenters and financial analysts. These sit one standard deviation below and above the mean WFH exposure, respectively. By 2025, our event-study estimates imply that, as compared to carpenters, roughly 1 in every 20 new hires for financial analyst roles has switched from junior to senior, and 1 in every 33 postings has switched to require more than 3 years of prior experience.

We do not interpret this evidence as ruling out strong impacts of GenAI on labor markets—our findings bear only on the relative hiring of junior and senior workers up to 2025. Nor do we seek to impugn the broader literature linking GenAI to declining demand for junior workers, much of which uses clever measurement strategies that do not rely on occupation-level exposure measures. We rather highlight that exposure-based designs can confound WFH with GenAI when no direct measure of either shock is available. Compared to the "AI-doomer" view, the fact that WFH plays a sizable role in the junior-hiring decline is, we think, cause for optimism. WFH delivers substantial benefits to workers and firms, and the organizational frictions it creates around supervising and developing early-career workers are surmountable through the diffusion of managerial practices better suited to remote-work environments. More dramatic policy remedies would be required if GenAI was the main driver of the junior-hiring decline.

The rest of the paper proceeds as follows. Section 2 reviews relevant literature, including studies that assess the role of remote working arrangements and organizational frictions that bias hiring away from junior employees. Section 3 formalizes how WFH can reduce the junior hiring share, using a stylized model of dynamic hiring decisions and worker-firm matching. Sections 4

describe the data and motivating facts around the fall of junior hiring intensity. Section 5 lays out our single- and two-treatment empirical designs and other aspects of our empirical strategies. Section 6 presents the main results, Section 7 summarizes additional results and robustness exercises, and Section 8 concludes.

2 Prior Studies of GenAI and WFH Effects

A fast-growing post-ChatGPT literature measures the labor market effects of generative AI. Many of these studies rely on empirical measures of the task-level exposure to these new technologies, which can be aggregated to occupations. Brynjolfsson et al. (2025a) use high-frequency US payroll data and find that workers aged 22–25 in the most exposed occupations experience a 16 percent relative employment decline after firm-level shocks are absorbed. Hosseini Maasoum and Lichtinger (2025) use US resume and vacancy data to identify firms posting for GenAI integrator roles and find a junior employment decline driven by slower hiring rather than separations or promotions. Teeselink (2025) estimates firm- and occupation-level difference-in-differences designs for the UK, and finds that highly exposed firms reduce employment, with the decline concentrated in junior positions. Azar et al. (2025) also use resume and vacancy data to estimate wage effects and find that average wages in exposed firms fall by 4.5 percent, while starting wages for junior and mid-level positions fall by 6.3 and 5.9 percent, respectively.

Many of these designs leverage exposure-based variation across tasks/occupations, which we argue raises two concerns. First, this variation can load heavily on the shift to WFH, so estimates based on GenAI exposure alone may pickup WFH effects. Second, they implicitly take the November 2022 release of ChatGPT as the forcing event for GenAI adoption, which may be too early in time to see labor market outcomes. The reorganization of work around WFH is therefore perhaps a more natural driver to associate with the change in hiring decisions that occurs after 2022. Moreover, the foundational assumption that GenAI will replace early-career workers is also contested. Althoff and Reichardt (2026) and Brynjolfsson et al. (2025b) emphasize the importance of GenAI as augmenting junior workers and reducing the barriers to entry across many jobs.

Studies which use survey-based adoption measures of GenAI also motivate our work. Humlum and Vestergaard (2025) link representative Danish technology adoption surveys to administrative labor records and estimate little effect of GenAI on earnings or hours worked two years after ChatGPT's release. Other survey evidence likewise finds widespread GenAI use but limited stated employment impact (Yotzov et al., 2026). Hartley et al. (2025) find similarly small labor market

effects despite widespread GenAI adoption. The gap between the findings of exposure-based GenAI effects and survey-level measures of actual GenAI adoption on employment may, in part, be explained by this paper's finding: exposure-designs risk misattributing GenAI effects to WFH effects.

The mechanism behind the effects of WFH proposed in this paper follow a long literature on how organizational frictions influence firms' hiring decisions. Firms hire early-career workers not only for current output, but for the future surplus created as they learn, benefiting from mentoring and managerial feedback.[2] Acemoglu and Pischke (1998) highlight explicitly this investment motive for early career hiring. This investment critically relies on the rate that workers learn on the job, and on the organizational structures that support this.[3] When effort is harder to monitor and output is harder to interpret, firms may be especially reluctant to hire early-career workers whose productivity, reliability, and work habits are not yet known (Holmström, 1982).

The evidence on the effects of WFH points directly to these organizational frictions. Yang et al. (2022) show that firm-wide remote work made collaboration networks more siloed. Emanuel et al. (2026) find that remote work has had scarring effects on young college graduates in remotable jobs, finding positive benefits of in-person office time. Atkin et al. (2023) show that home-based work lowered productivity partly due to slowing learning. Aksoy et al. (2025) show that in-person onboarding raised later productivity and reduced attrition, even in a fully remote firm. Aksoy et al. (2026) further found that just one office day per month, relative to fully-remote work, raised productivity and increased communication between workers and their managers. Emanuel et al. (2026) and Emanuel et al. (2024) show that proximity to co-workers increases feedback among female software engineers, with the largest gains for younger workers.

These organizational frictions associated with WFH must be balanced by the wide-ranging benefits.[4] Moreover, workers highly value WFH arrangements (Mas and Pallais, 2017; Cullen et al., 2025). A reversion to pre-pandemic work arrangements is therefore unlikely. To address the monitoring, supervision, and learning challenges that WFH arrangements pose will require the diffusion of new management practices that better accommodate hybrid and remote workers in the early-stages of their career.

[2] This is the standard logic of learning-by-doing and human-capital accumulation; see Arrow (1962), Becker (1964), Mincer (1974).

[3] On supervision, knowledge hierarchies, peer learning, and proximity, see Garicano (2000), Prendergast (1999), Lazear and Oyer (2013), Lazear et al. (2015), Sen (2024), Battiston et al. (2021), and Jarosch et al. (2021).

[4] For example, documented benefits of WFH work arrangements including higher labor force participation (e.g. Bloom et al. (2026)), increased productivity and staff retention (e.g. Bloom et al. (2015)), new business creation (e.g. Kwan et al. (2025)), and location flexibility (e.g. Akan et al. (2025) and Lambert and Larkin (2024)).

Before we turn to our main empirical findings, the next section presents a minimal conceptual framework that explains the organizational frictions associated with WFH in a model with two-tiered employment, supervision and on-the-job learning, and dynamic hiring decisions.

3 Conceptual Framework

This section presents a partial-equilibrium model in which firms hire senior and juniors workers into teams that are subject to supervision constraints and in which junior workers engage in on-the-job learning. The model is purposefully simple, with the main aim of illustrating conceptually how changes to the rate of on-the-job learning and supervision burden, both engendered by the reorganization of work around WFH, impact the relative value of senior and junior workers within the firm.[5]

A risk-neutral firm employs juniors (J) and seniors (S) in a team, in quantities ℓ_J and ℓ_S. Each junior uses a fraction $0 < \phi < 1$ of senior time and supplies $0 < \eta < 1$ efficiency units of junior labor.[6] Firms take wages $w_J < w_S$ as given. Effective senior and junior time in production, net of organizational frictions, are $s \equiv \ell_S - \phi\ell_J$ and $j \equiv \eta\ell_J$. Flow profit is

$$\pi = F(s, j) - w_S\ell_S - w_J\ell_J, \tag{1}$$

where F aggregates effective worker time and has standard properties.[7] We assume an interior solution throughout, with both worker types employed. It is useful to introduce the junior-to-senior productivity wedge

$$\psi \equiv \frac{\eta F_j - \phi F_s}{F_s}, \tag{2}$$

so that an additional junior contributes $\psi F_s = \eta F_j - \phi F_s$ to static flow output.

We now turn to the hiring dynamics. Let V_S and V_J denote the firm's continuation values from a marginal senior and junior worker. Juniors become seniors through on-the-job learning at rate

[5] Proofs of all main results are in Appendix A. A general-equilibrium extension with endogenous wages is in Appendix B.

[6] This reduced-form formulation is motivated by models of hierarchical production in which lower-level workers rely on scarce higher-level time for monitoring or problem solving. In monitoring models, supervisors' limited time constrains span of control (Calvo and Wellisz, 1978; Qian, 1994). In knowledge-hierarchy models, production workers rely on higher-level problem solvers when they encounter problems they cannot solve (Garicano, 2000; Garicano and Rossi-Hansberg, 2006).

[7] Formally, F is twice continuously differentiable, strictly increasing, concave, homogeneous of degree one, strictly concave in the factor ratio, satisfies the Inada conditions, and has finite positive local elasticity of substitution.

$\lambda > 0$, and workers separate at rates δ_J, δ_S.[8] Hiring occurs in separate labor markets with seniority-specific convex recruitment costs $C_i(h_i) = \frac{1}{2}\kappa_i h_i^2$, $i \in \{J, S\}$.[9] In steady state, continuation values for each worker satisfy:

$$V_S = \frac{F_s - w_S}{r + \delta_S}, \qquad V_J = \frac{\psi F_s - w_J + \lambda V_S}{r + \lambda + \delta_J}. \tag{3}$$

where F_j and F_s denote partial derivatives and r is the discount rate. V_S is standard. V_J reflects both the organizational environment as well as the investment channel.

At any point in time, firms hire workers up to the point where the value of a marginal worker is equal to the (linear) marginal cost of hiring. This implies an optimal junior share of hiring:

$$z^\star \equiv \frac{h_J}{h_J + h_S} = \left[1 + \frac{\kappa_J}{\kappa_S}\frac{V_S}{V_J}\right]^{-1} \tag{4}$$

We next present comparative statics to understand the mechanism behind WFH and junior-hiring. We model WFH as raising supervision costs, ϕ, and lowering the promotion rate, λ.[10] These comparative statics are taken at a fixed team composition: the current worker stocks ℓ_S and ℓ_J are held at their existing levels when the parameter changes, so only the mechanical dependence of effective senior and junior time ($s \equiv \ell_S - \phi\ell_J$ and $j \equiv \eta\ell_J$) on ϕ and η feeds through to marginal products.[11] The relevant comparative statics are:

$$\text{Learning rate:} \quad \frac{\partial z^\star}{\partial \lambda} > 0, \tag{5}$$

$$\text{Supervision cost:} \quad \frac{\partial z^\star}{\partial \phi} < 0, \tag{6}$$

both of which hold as long as

$$\frac{\lambda}{\lambda + r + \delta_J} < \frac{V_J}{V_S} < 1. \tag{7}$$

[8] The transition rate λ is a reduced-form representation of promotion from junior to senior status. Such progression can reflect job assignment (Sattinger, 1975; Rosen, 1982; Waldman, 1984), on-the-job human-capital accumulation (Becker, 1964; Ben-Porath, 1967; Hashimoto, 1981), or employer learning about worker ability (Harris and Holmström, 1982; Holmström, 1999; Farber and Gibbons, 1996). Gibbons and Waldman (1999) integrate these mechanisms in a theory of wage and promotion dynamics inside firms.

[9] The quadratic specification follows a common assumption about labor adjustment costs (see e.g. Sargent, 1978) and is consistent with evidence of diseconomies of scale in recruitment and increasing marginal hiring costs (Manning, 2006; Pfann and Verspagen, 1989; Blatter et al., 2012)

[10] Section 2 discusses evidence for both mechanisms.

[11] Formally, the derivatives below are partial derivatives of the continuation values V_S and V_J and of the hiring share $z^\star$ at a given (ℓ_S, ℓ_J), rather than total derivatives that allow the stocks themselves to re-equilibrate. They tell us how the firm's instantaneous hiring choice tilts toward or away from juniors at the moment WFH shifts ϕ and λ, holding the current team fixed. The induced stock-flow adjustment in ℓ_S and ℓ_J plays out subsequently through the laws of motion governed by λ, δ_J, and δ_S and the new hiring flows h_J, h_S. Appendix A states the fixed-stock convention formally and shows that the sign predictions are preserved once wages clear endogenously.

The first comparative static in 5 highlights that a slower learning rate lowers the junior-hiring share, and holds when the upper bound of Inequality 7 is satisfied. This condition is met when promotion delivers a capital gain to the firm, i.e. $V_S > V_J$.[12] Comparative static 6 states that greater supervision requirements further reduce the junior hiring share, and holds when the lower-bound condition of Inequality 7 is satisfied. When this is the case, so that $V_J > \frac{\lambda}{\lambda+r+\delta_J} V_S$, a junior worker's value exceeds the discounted capital gain from promotion alone. This holds when juniors retain at least as much wage surplus relative to their marginal product as seniors do, i.e. when $w_J \leq \psi w_S$.[13]

Appendix A gives the formal proofs and we focus here on the intuition for the direction of each effect. A lower learning rate λ makes juniors less attractive as an investment: the capital gain from promotion is more heavily discounted, so steady-state junior hiring falls. A higher supervision cost ϕ operates through a twofold channel: it both reduces the marginal product of a junior worker, equal to $\psi F_s = \eta F_j - \phi F_s$ by the wedge definition in (2), since each marginal junior produces ηF_j but draws ϕF_s of senior supervisory time, and raises the marginal productivity of senior time F_s since $s \equiv \ell_S - \phi \ell_J$.

In this partial-equilibrium framework, we have taken wages as fixed. Appendix B extends the model by letting the senior and junior wages adjust to clear type-specific hiring markets. The sign of both comparative statics is preserved. Endogenous wages only dampen the magnitude of the quantity response, because part of the WFH-induced shift in hiring demand is absorbed by wages rather than by quantities.[14]. Having provided conceptual evidence of the channels via which WFH may impact the relative value of junior and senior workers within the firm, we next turn to a discussion of the data.

[12] Equivalently, $(F_s - w_S)/(\psi F_s - w_J) > (r+\delta_S)/(r+\delta_J)$. This condition is easier to satisfy when seniors are more attached to the firm than juniors. Empirically, job mobility is concentrated early in workers' careers and declines with age, experience, and tenure (Topel and Ward, 1992). Promotion can therefore generate a capital gain even when senior per-period surplus is somewhat smaller than junior per-period surplus. See Appendix A for the full derivation.

[13] Since $V_S > 0$ implies $F_s > w_S$, this wage-wedge condition delivers the simpler product-side inequality $\psi F_s > w_J$: a junior's static marginal product net of supervision cost exceeds their wage. Both typically hold: juniors accept a wage markdown in return for the amenity value of on-the-job training, reputational/signaling gains, and access to the firm's internal job-ladder (Barron et al., 1989; Pallais, 2014; Stanton and Thomas, 2016; Gibbons and Waldman, 1999).

[14] The extension also delivers wage predictions: higher ϕ raises the senior wage and lowers the junior wage whenever the static-productivity loss from supervision exceeds the attenuated continuation-value gain from promotion, and a lower λ lowers the junior wage whenever junior hiring is an investment in future senior surplus ($V_S > V_J$). Empirically, senior-wage premia moves in ways consistent with these GE predictions (see Appendix E)

4 Data and Motivating Facts

This section describes the two microdata sources we use to measure hiring and recruitment, as well as our measures of occupation-level exposure to GenAI and WFH. We also close this section with a discussion of our two main motivating facts, which are the fall in the share of all new hires going to early-career workers, and the strong overlap between measures of GenAI and WFH exposure used in the literature.

4.1 Data Sources and Variable Construction

Throughout this paper we rely on two microdata sources to measure worker hiring and job vacancy postings. Both sources cover the United States, the United Kingdom, Canada, and Australia over 2017–2025. We aggregate these data into two panels at the state-by-6-digit ONET-SOC occupation code-by-month level, and at the firm-by-2-digit occupation-by-year level.[15] We use the terms 'broad occupation' to refer to a 2-digit ONET-SOC occupation group, and 'detailed occupation' to refer to 6-digit O*NET-SOC occupation.

Data on new hires comes from Revelio Labs, which provides monthly matched employer–employee records assembled from résumés (predominantly, but not exclusively, from LinkedIn). Each record contains the employer, occupation, start date, and location. The data provider further classifies seniority into seven categories, parsed from the job title and the worker's prior employment history. Our definition of 'junior' combines the bottom two groups from these categories.[16] Sampling weights, which adjust for differential coverage across occupations and geographies, are applied throughout.[17] Our main outcome of interest used in our analysis is the 'junior' share of all new hires in a given cell.

Our measure of online job vacancies comes from Lightcast, who supplies the near-universe of online job vacancy postings collected from thousands of online job boards, official recruitment websites, and employers' websites. These data include occupation, employer name, region or city, and the number of years of relevant experience the employer requires of applicants. Missingness in the years of experience field is non-trivial, and the absence doesn't always imply zero years experience are required. To address this, we apply an inverse-probability weighting

[15] In the UK, we treat the devolved nations of England, Wales, Scotland, and Northern Ireland as our 'State'.

[16] Revelio Labs assigns each position to one of seven levels ('Entry-Level', 'Junior', 'Associate', 'Senior', 'Manager', 'Director', 'Executive').

[17] The data provider makes efforts to ensure representativeness, so as to address the potential skew towards white-collar jobs from online résumé sources. See further remarks about the representativeness of these data in Appendix C

(IPW) correction prior to aggregating. Additional details on this are in Appendix C.2. Our main outcome of interest from these data is the share of all postings in a given cell that require three or fewer years of experience.

We measure *exposure* to generative AI and WFH at the occupation level using two main data products. GenAI exposure is measured across narrow occupations using the index of Eloundou et al. (2024). These authors use a language model to assess which tasks can be completed in at least 50% less time with the use of generative AI tools, and aggregate this to the occupation-level based on the fraction of tasks in an occupation that meet this criterion. This has been widely used in the literature. Our measure of WFH exposure comes from Hansen et al. (2023), who classify job postings which offer one or more days of remote/hybrid work from online job postings. The share of postings which meet this criteria in 2021-22 in each occupation is built using the entire four-country sample, and we use this to define occupation-level exposure to WFH. Both our GenAI and WFH exposure measures vary only across occupations. We z-score standardize both measures across occupations, after removing five 2-digit O*NET groups and residual/'catch-all' categories.[18] In robustness exercises, we also use the teleworkability index of Dingel and Neiman (2020), and two further AI-exposure indices from Felten et al. (2023) and Appel et al. (2025). When used, these alternative measures are also z-score standardized across the same set of occupations.

Finally, we also utilize a measure of *actual WFH adoption*. This again comes from Hansen et al. (2023), but this time we use postings from specific firms or regions to see how intensely WFH arrangements were offered in these focal units in 2021-22. This allows us to identify which units were explicitly offering WFH arrangements prior to our analysis period of 2023-25. We are able to link this job posting based measure to our hiring outcomes by matching employer-names across our two data sources (see Appendix C.3 for the matching procedure and full performance metrics.) A symmetric measure of actual GenAI adoption would be a natural complement, but is presently not available at the scale and coverage of our WFH measure.

4.2 Motivating Facts

The decline in hiring and recruitment of less experienced workers has been widely reported in the news (Almeida, 2025; Shrivastava, 2025; Raman, 2025; Thompson, 2025), and the underlying

[18] The excluded occupations in our analysis are the two-digit categories 21, 25, 33, 45, 55, covering community/social services, education, protective services, farming/fishing/forestry, and military occupations. Hiring in these occupations is often centralized, lacks meaningful variation, or else doesn't exist in both exposure measures. The catch-all occupations (e.g. '51-9199 Production Workers, All Other') are not used as they are unavailable in the Eloundou et al. (2024) data.

facts are also presented across multiple data sources, including payroll data (see Brynjolfsson et al. (2025a)) and hiring data (see Hosseini Maasoum and Lichtinger (2025)). As these facts have been widely reported, we offer just a few additional exhibits constructed from our own sample:

First, we show in Figure 1 the extent to which the junior share of new hires has declined across all four countries since late 2022. Each country's series is stable through 2019, dips during 2020–2021, and then breaks downward from late 2022, reaching roughly 8 to 11 percentage points below the 2019 baseline.[19] The synchronicity across four economies with different institutions, immigration regimes, and macro trajectories is the strongest visual argument for a common underlying shock.

Second, we highlight that our two candidate exposure measures of WFH and GenAI strike largely the same occupations. Figure 2 shows this correlation, depicting the rank of 683 6-digit O*NET-SOC occupations. The rank correlation between our WFH and GenAI exposure measures across all occupations is $\hat{\rho}_S \approx 0.77$. Software developers, accountants, and management consultants sit at the top of both rankings; electricians, janitors, and construction laborers at the bottom. This makes plain the omitted variable problem that may surface when either shock is analyzed in isolation of the other. We next turn to our main empirical strategy, which is designed to disentangle the relative role of WFH and GenAI in the decline of early-career hiring.

5 Empirical Strategy

We exploit the timing of the rise of WFH and GenAI, using event-study/difference-in-differences empirical designs. The material absence of mass adoption of WFH and GenAI in our 2017-2019 pre-period serves as a natural baseline. The main post-period of interest is 2023-2025. In our main specifications, the identifying variation comes from the differential exposure of units to WFH and GenAI shocks, based either directly on the narrow occupation or on the pre-period occupation mix in a given unit. Our more direct designs which use actual WFH adoption in 2021-22 instead exploit variation based on observed WFH usage.

We use the following notation throughout: n indexes firms, b is a 'broad' 2-digit occupation group, j is a 'detailed' 6-digit O*NET-SOC occupation, m is a state, and c is a country. We further write $J(b)$ to denote the set of detailed occupations contained in broad group b. A regression

[19] Across the post-period, entry-level hires are down 14 to 29 percent across countries and junior hires are down 12 to 26 percent, while senior hires are 5 to 21 percent *above* the pre-period average. Appendix Figure D.2 decomposes the US decline into absolute flows by seniority level and Appendix Table C.3 reports the cross-country averages.

unit i is either a state×detailed-occupation cell (m, j) or a firm×broad-occupation cell (n, b). Our exposure measures are given by Z_i^k, where k indexes the shock of interest (WFH or GenAI).

Using our data on new hires and job postings, the outcomes of interest are all share measures, expressed in percentage points. We construct these shares using two different aggregations. Our first design constructs outcomes and exposures using cells defined at the state×detailed-occupation units at the monthly frequency. Our second design uses cells defined at the firm×broad-occupation unit measured at an annual frequency. The former level of aggregation exploits within occupation-state/region time-series variation, while the latter exploits within-firm variation.[20] Both designs further allow us to absorb a rich set of controls and suitably define treatment either directly from our measures (using narrow occupations) or else using a shift-share design based on a given unit's pre-period occupation mix. We use the unit i notation to refer to both designs interchangeably.

5.1 Exposure-based Designs

Our treatment (Z) is the standardized exposure to WFH and GenAI exposure. Let y_{it} denote the outcome of interest for unit i and period t. Our main event-study specification is therefore:

$$y_{it} = \sum_{k \in K} \sum_{\tau \neq 2019} \phi_\tau^k \left(Z_i^k \times \mathbf{1}\{\text{year}(t) = \tau\} \right) + \alpha_i + \delta_t + \varepsilon_{it}, \tag{8}$$

where shock k belongs to $K = \{\text{WFH}, \text{GenAI}\}$. With 2019 omitted, ϕ_τ^k is our estimate of interest and captures the year-specific exposure gradient associated with a one-standard-deviation increase in the relevant exposure measure, Z_i^k, relative to the omitted period. For compact tables we estimate the corresponding pre/post specification with $\text{Post}_t \equiv \mathbf{1}\{\text{year}(t) \geq 2023\}$, so that the estimating equation is:

$$y_{it} = \sum_{k \in K} \beta^k \left(Z_i^k \times \text{Post}_t \right) + \alpha_i + \delta_t + \varepsilon_{it}. \tag{9}$$

In the above two-period design, β^k captures the post-period exposure gradient associated with a one-standard-deviation increase in exposure to either WFH or GenAI.[21]

[20] Appendix E.6 shows the density plots of the pre-control variation at both levels of aggregation, along with the residual plots after absorbing unit and time fixed effects.

[21] With heterogeneous effects, β^k averages unit-specific responses, weighted by the regression weights and by the residual exposure variation that survives partialling out fixed effects, controls, and the co-exposure. We allow the two shocks to interact in additional specifications by adding $Z_i^{\text{WFH}} \times Z_i^{\text{GenAI}} \times \text{Post}_t$ to (9) and the corresponding year-specific triple interactions to (8).

When we aggregate data into firm-by-broad occupation units, the exposure measures of GenAI and WFH are constructed as a weighted shift-share measure across a units 2017-2019 narrow-occupation mix of hiring/postings. With these shares given as $s^{\text{pre}}_{nb,j}$, the shock is then given by $Z^k_{nb} = \sum_{j \in J(b)} s^{\text{pre}}_{nb,j} Z^k_j$.

We use the same granular fixed effects and inference strategy for both data-sets, but these differ depending on the level of aggregation. In the state×detailed-occupation panel, α_i is a state×detailed-occupation unit fixed effect and δ_t is a state×month time fixed effect. Standard errors are clustered by both detailed occupation and state in these designs. In the firm×broad-occupation panel, α_i is a firm×broad-occupation fixed effect and δ_t is a country×year fixed effect.

Since we are using a shift-share exposure measure when we work at the firm-by-broad-occupation unit level, we report shock-level standard errors following Borusyak et al. (2022). Our estimation procedure uses long-differences of the outcome within each unit, which absorbs the firm×broad-occupation fixed effect. We next aggregate the long-differenced outcome to the 6-digit-occupation shock level using each unit's pre-period occupation-mix as weights. We then partial out country fixed effects at the shock level and clusters standard errors at the broad-occupation (2-digit) aggregation of the shock identifier. This approach follows Borusyak et al. (2022), but is asymptotically equivalent to Adão et al. (2019).

Identification rests on an exposure-parallel-trends condition. Because Z^{WFH}_i and Z^{GenAI}_i are continuous exposure measures, Equations (8) and (9) are continuous-treatment DiD designs in the sense of Callaway et al. (2025) and De Chaisemartin et al. (2024). Our maintained identifying condition is that, after partialling out the included fixed effects and, in the joint design, the co-exposure, units with different pre-period WFH or GenAI exposure would not have experienced differential post-2022 breaks in early-career hiring absent the relevant shock.[22] This condition supports interpreting ϕ^k_τ and β^k as reduced-form exposure-gradient DiD coefficients.Interpreting these coefficients as average causal responses to a one standard deviation increase in WFH/GenAI exposure further requires a stronger no-selection-across-dose condition (Callaway et al., 2025).

[22] Formally, for focal shock $k \in \{\text{WFH}, \text{GenAI}\}$ and co-exposure $\ell \neq k$, $E[Y_{it}(0) - Y_{i,2019}(0) \mid Z^k_i, Z^\ell_i, \mathscr{C}_{it}]$ is constant in Z^k_i, where $\mathscr{C}_{it}$ denotes the fixed-effect structure and other controls partialled out in the relevant specification.

5.2 WFH Adoption Designs

As a final main exercise, we also analyze the effects of *actual WFH adoption* during 2021–2022 as treatment. At the firm-by-broad occupation level, we use a binary indicator for whether the unit saw any postings offering WFH, D_{nb}^{WFH}, since continuous shares are too sparse at this cell to identify a dose-response. At the state-by-detailed occupation level, we use a standardized measure of the share of postings offering WFH in that state-occupation cell during 2021-22, D_{mj}^{WFH}. We then estimate:

$$y_{nbct} = \gamma\left(D_{nb}^{\text{WFH}} \times \text{Post}_t\right) + \alpha_{nb} + \delta_{ct} + u_{nbct}, \tag{10}$$

$$y_{mjct} = \gamma\left(D_{mj}^{\text{WFH}} \times \text{Post}_t\right) + \alpha_{mj} + \delta_{ct} + u_{mjct}, \tag{11}$$

with firm×broad-occupation and country×year fixed effects in (10) and state×detailed-occupation and country×year fixed effects in (11).

Because WFH adopters may differ from non-adopters on observable exposure, we employ a matching strategy in the estimation of (10) and (11). This matches units based on their exposure to both WFH and GenAI using coarsened exact matching (CEM) on a 5×5 grid of WFH and GenAI exposure quintiles (Iacus et al., 2012; Ho et al., 2007).[23]

All joint specifications report variance-inflation factors (VIF) and a condition number (CN) for the focal treatment block, computed from the residualized, weighted treatment columns rather than from the raw exposure indices, together with partial F-statistics for each treatment. We use these less conventionally reported statistics to diagnose the potential issues of multi-collinearity, and report the methods used to produce them in Appendix F.2. We now turn to the main results related to the empirical designs described above.

6 Main Results

This section presents the empirical evidence in two steps. First, we estimate event-study and diff-in-diff specifications that include WFH exposure and GenAI exposure both separately and jointly, at two levels of aggregation and on two outcomes. Second, we move from predetermined exposure measures to using *actual* WFH adoption in 2021–2022 as our treatment. Across these

[23] Matching covariates are $Z_{nb}^{\text{WFH}}, Z_{nb}^{\text{GenAI}}$ in the firm design and $Z_j^{\text{WFH}}, Z_j^{\text{GenAI}}$ in the occupation design. Strata without within-stratum treatment variation are dropped; regression weights equal postings or hires multiplied by the CEM weight.

designs, the WFH-channel remains a robust predictor of a fall in junior-share of new hires and share of vacancy postings requiring no more than three years experience. By contrast, the GenAI exposure effects attenuate sharply in joint-designs.

6.1 Exposure-based WFH and GenAI Effects

When entered separately, the WFH and GenAI exposures generate near-identical event-study paths. Entered jointly, only the WFH path remains stable. Figure 3 shows the event-study coefficients from Equation 8 estimated on firm×broad-occupation×year cells, with each panel depicting both a single-treatment path (black circles) and the corresponding joint-treatment path (colored diamonds). The black single-treatment paths are nearly indistinguishable across the two exposures, with each predicting a 2–3pp post-2022 decline in the junior share of new hires and around a 2pp decline in the share of postings which require no more than 3 years of experience. We see no evidence of pre-trends. Once both exposures enter jointly the WFH path (blue diamonds, panels (a) and (c)) remains essentially unchanged compared to the single-treatment coefficient path, while the measured-GenAI path (red diamonds, panels (b) and (d)) attenuates heavily and becomes statistically insignificant.

These patterns also hold when we estimate these coefficient-paths using a different level of aggregation. The same design as Equation 8, estimated on state×detailed-occupation×month cells is shown in Figure 4, where we again see that the WFH coefficient path is invariant to the inclusion of the GenAI exposure co-treatment, and the GenAI exposure path again is heavily attenuated and unstable when WFH exposure co-treatment is added. Pre-trends in the postings designs are not flat in this panel, but if anything they move in the opposite direction to the main effects. The postings decline begins somewhat earlier, around 2020, rather than 2023. The timing is consistent with the 2020–21 WFH adoption shock binding into hiring strategy with a lag, as emergency WFH arrangements crystallised into a more permanent amenity over 2022–23.

Moving away from event-studies, Table 1 reports the corresponding two-period difference-in-differences estimates from Equation 9, allowing a clean comparison of the single-treatment (cols (1),(2),(5),(6)), additive (cols (3),(7)), and interaction designs (cols (4),(8)) across both panels and outcomes. Comparing columns (1) and (3), the WFH exposure coefficient on the junior share of new hires falls only modestly—from −1.69 to −1.42 pp in the state×occupation panel and from −1.88 to −1.57 pp in the firm×occupation panel—when measured-GenAI exposure is added jointly. The GenAI exposure coefficient falls from −1.34 to −0.41 pp and from −1.59 to −0.45 pp when WFH is added, and on the postings outcome it changes its sign. This asymmetry points

to measured WFH exposure as the more salient shock in these designs. The diagnostics show that the asymmetric attenuation is not driven by inflated variance from collinearity. Variance-inflation factors stay below 2.4 and condition numbers below 2.7 across both panels.

To put these findings on a concrete scale, consider occupations roughly one standard deviation below the mean WFH-exposure—skilled trades such as construction carpenters and auto-mechanics—and those who lie roughly one standard deviation above the mean, in white-collar analytic roles such as financial analysts and information-systems managers. Reading off the 2025 coefficients in Figure 3, a 2-standard-deviation gap translates into around 4–5pp larger decline in the junior-hire share and a roughly 3pp larger decline in the share of postings requiring no more than 3 years of experience—equivalent to shifting up to 1 in every 20 new hires from junior to experienced workers, and raising the experience bar above 3 years on up to 1 in every 33 postings, relative to the blue-collar occupations.

6.2 Actual WFH Adoption

We now report the results from our second main exercise, namely the use of data which measure which units *visibly adopted* remote or hybrid work in 2021–2022. We find that units where visible adoption of WFH was strong subsequently shifted away from early-career hiring. This is estimated after matching on pre-period WFH and GenAI exposure. The treatment variables are D_{nb}^{WFH} and D_{mj}^{WFH}, as defined in Section 4. Event-study paths in Figure 5, the year-by-year analogues of Equations 10 and 11, show flat pre-trends through 2019 in three of the four panels and sharp post-2022 declines.[24] The pre/post estimates in Table 2 (Equations 10 and 11) match this picture, with the firm-level binary actual-WFH coefficient between −0.69 and −1.20 pp ($p < 0.01$) in every column of Panel A, and scarcely smaller in the joint specifications than in the single-treatment ones. The occupation-level continuous actual-WFH coefficient attenuates when paired with the WFH exposure index (Panel B), as expected since both are measured at the same level and identify off largely the same variation.

Overall our estimates suggest that the shift towards WFH arrangements which was forced during the pandemic, and which persisted thereafter, is a robust predictor of declining relative demand for early-career workers. The variation in GenAI-exposure that produces similar effects in single-treatment designs appears to not remain once the WFH co-treatment is included. In

[24] The exception is panel (b), the firm-level postings outcome, where the 2017–2019 series drifts by roughly 0.20 pp per year. The post-2022 trajectory deviates only modestly from the extrapolated pre-trend in this panel, so the panel (b) point estimates should be read as an upper bound on the treatment-attributable decline. The other three panels show no comparable drift.

our next section, we briefly discuss the host of robustness exercises and checks, along with additional results. Across alternative specifications, the use of non-linear co-treatment controls, and exercises that tackle the potential role of measurement error, the consistent pattern is that WFH remains a robust predictor of the fall in early-career hiring.

7 Additional Results and Robustness

The main result in this paper is that WFH, measured both using occupational exposure and actual adoption, predicts a sizable decline in the junior-share of new hires, and the share of job postings requiring no more than three years experience. The goal of this section is to briefly explain all the additional results and robustness checks we have deployed, and explain why we interpret these as favorable to this main conclusion. This includes using alternative measures of GenAI and WFH exposure, testing sensitivity to functional form and linear OLS models, assess if the main result is overly sensitive to any one country or occupation group, and assess whether measurement error or unobserved confounding factors could plausibly overturn the main result. This section offers brief remarks, with more detailed descriptions of each item in the Appendix.

Auxiliary wage and replacement margins. Two further margins implied by our framework move with WFH exposure in the predicted direction, and not with measured GenAI exposure. If WFH raises supervision cost or lowers the learning rate (Section 3), the senior–junior starting-wage gap should rise and the junior-to-senior replacement rate should fall. Appendix Table E.1 estimates the joint specification on both auxiliary outcomes at the state×occupation and firm×occupation levels. Both margins move with WFH exposure in the predicted direction and across both panels, while the GenAI coefficient runs the wrong way on each margin—statistically significantly so in the firm-level panel.

Leave-one-out occupational robustness. The WFH-dominant ranking is not identified off any single 2-digit occupational cluster. The concern is that a particular major group (for example, Computer and Mathematical occupations, which combine high remote-work intensity with the mass tech layoffs that began in late 2022) drives the result on its own. Appendix Table E.3 drops each of the eighteen included 2-digit ONET-SOC major groups in turn and re-estimates the joint specification. The WFH coefficient remains negative and statistically significant in every leave-one-out replication across all four designs. The largest movement occurs when Office and Administrative Support (SOC 43) is dropped, attenuating the WFH coefficient on new hires from -1.42 to -1.11 at the state×occupation level and from -1.57 to -1.14 at the firm level, with both remaining significant at conventional levels.

Country-by-country robustness. The result is also not identified off any single country. Appendix Table E.2 re-estimates the joint specification separately for the US, UK, Australia, and Canada. On the new-hires outcome, the WFH coefficient is negative and statistically significant in all four countries at both aggregation levels, although in UK new hires at the state-occ level GenAI is also negative and significant—the one country×design cell where both exposures survive jointly. On the postings outcome the effect holds in the US and Australia at both levels and in Canada at the firm level, with the statistically null WFH estimates confined to both UK postings panels and the Canadian occupation-level postings panel. Across all sixteen country×design combinations the WFH-dominant pattern is preserved in thirteen.

Non-parametric co-treatment controls. The WFH-dominant ranking is not an artifact of imposing a linear control on GenAI exposure. A linear control could in principle miss non-linear dependence between WFH and GenAI exposure, leaving residual GenAI variation that the WFH coefficient then absorbs. Appendix Figure F.1 therefore replaces the linear control with step functions of progressively finer quantile resolution, ranging from no control through deciles. The WFH coefficient is stable at every resolution. The GenAI coefficient, by contrast, falls sharply once even a single above-median WFH dummy enters as the co-treatment and is statistically indistinguishable from zero at every finer resolution—on the postings outcome it often flips slightly positive in sign.

Selection on unobservables. An unobserved confounder would have to be substantially stronger to overturn the WFH coefficient than the GenAI coefficient. The threat is that an unmeasured occupational characteristic correlated with both exposures accounts for the residual joint coefficients. Appendix Table F.1 probes this using the Cinelli and Hazlett (2020) robustness value $RV_{q=1}$ for the single-treatment design of each exposure (in the tradition of Altonji et al., 2008; Oster, 2019). The Cinelli–Hazlett $RV_{q=1}$ is 3–5 times larger for WFH than for GenAI (around 0.05–0.09 vs. 0.01–0.03): a confounder explaining only 1–3% of the residual variation in treatment and outcome would suffice to drive the GenAI coefficient to zero, against 5–9% for WFH.[25]

Measurement-error tests using Monte Carlo simulation. In Appendix Table F.2, we run simulations to assess whether a true GenAI exposure effect would spuriously show up as a WFH exposure effect, when the former is noisy and the latter is not. We do this by first proposing a data generating process where the full effect is loaded on to GenAI, and WFH is mechanically correlated but has no true underlying signal. We then introduce Gaussian noise to the GenAI exposure measure, and run our main estimates 500 times at different noise-to-signal ratios (ν).

[25] The same diagnostic applied to the GenAI coefficient *in isolation* — without WFH in the regression — would itself have flagged the GenAI result as not robust to plausible unobserved confounding.

We then estimate our two-period DiD estimation using 500 replication draws for each level of ν. Table F.2 Panel A shows we would need implausibly large noise ($\nu = 1$) to cross the threshold where, in 5% of draws or more, we would attribute the true effect of GenAI on the junior-share of hires to WFH. For the share of job postings requiring no more than 3 years experience, the misattribution rate is zero across all ratios, even implausibly large ones. This exercise suggests that classical measurement error in the GenAI exposure measure is very unlikely to explain the WFH-dominant ranking.

Robustness across alternative WFH and GenAI exposure measures. The WFH coefficient remains negative and statistically significant under every alternative-measure pairing. The potential threat is that the result depends on the specific Hansen et al. (2023) WFH measure and the Eloundou et al. (2024) GenAI measure that anchor our main specifications. Appendix Table E.4 therefore re-estimates the joint specification with the WFH×GenAI interaction (columns (4) and (8) of Table 1) across all six combinations of two WFH measures (Hansen et al. (2023) and the Dingel and Neiman (2020) teleworkability index) and three GenAI measures (Eloundou et al. (2024), Felten et al. (2023), and the Anthropic Economic Index of Appel et al. (2025)), with missing occupations imputed by hierarchical 5/3/2-digit ONET means. The WFH coefficient is larger in magnitude than the GenAI coefficient in 22 of 24 cells. Both exceptions arise in the firm-level junior-hires panel under the Dingel and Neiman (2020) WFH measure: the Felten et al. (2023) GenAI coefficient is larger in magnitude (-1.57 vs. -0.89), and the Eloundou et al. (2024) coefficient is marginally larger (-1.23 vs. -1.22).

Residualized single-treatment designs. Partialling each exposure out of the other before estimation preserves the WFH-dominant ranking. This is a simple visual check on the joint design that avoids any concern about controlling for a collinear regressor. Appendix Table E.5 reports residualized single-treatment regressions in which each exposure is partialled out of the other at the occupation level. The WFH-residualized-on-GenAI series remains a robust negative predictor of post-2022 declines in early-career hiring. The GenAI-residualized-on-WFH series, by contrast, is small, statistically insignificant, or wrong-signed.

Obviously-related instrumental variables. Correcting for classical measurement error in the GenAI index using obviously-related IV largely preserves the WFH-dominant ranking. The threat is again that the GenAI coefficient is attenuated by classical noise in its index, but here we use the multiple available GenAI indices as instruments for one another rather than simulating the noise. Appendix Table F.3 applies the obviously-related-IV correction of Gillen et al. (2019) to the three GenAI measures, instrumenting each with the mean of the other two. The WFH coefficient remains negative and significant in five of six rows. The lone exception (Panel A under the

Appel et al. (2025) Anthropic Economic Index) returns an implausibly negative implied reliability $\hat{R}$—part of a broader pattern in which $\hat{R}$ ranges from -11 to $+5$ across the three measures, itself inconsistent with a single classical-measurement-error story under which $\hat{R}$ would be common and lie in $(0,1)$.

Alternative specifications and samples. The WFH-dominant ranking is robust to the experience threshold used to define "junior" (Appendix E.6 varies the threshold from ≤ 1 to ≤ 6 years), the choice of estimator (Table E.7 re-estimates under unweighted OLS, the fractional Poisson pseudo-likelihood of Papke and Wooldridge (1996), and OLS on the log-odds transform), four-country pooling (Appendix E.2), and the IPW correction for unreported experience in the postings data (Appendix C.2). The lone exception is the occupation-level new-hires outcome under the log-odds transform, where both WFH and GenAI are statistically insignificant.

In sum, the WFH-dominant ranking—a stable, negative, statistically significant joint coefficient on WFH exposure alongside an attenuated and frequently insignificant joint coefficient on measured GenAI exposure—survives every robustness exercise reported here: alternative co-treatment functional forms, leave-one-out occupational and country variation, selection-on-unobservables diagnostics, classical measurement-error simulations, the obviously-related-IV correction, alternative WFH and GenAI exposure indices, and alternative outcome definitions, estimators, and samples. The same is not true for the measured-GenAI coefficient, which attenuates, flips sign, or becomes statistically insignificant under most of the same exercises.

8 Conclusion

Demand for junior talent appears to have fallen in the post-pandemic era. This paper tests two possible explanations for this shift: the impact of generative AI tools in the workplace, and the rapid and persistent adoption of work-from-home arrangements. Separately identifying the effects of both is a challenge, because widely used occupation-level measures of WFH and GenAI exposure are highly correlated, so single-treatment designs risk misattributing one channel to the other.

Using two microdata sources, across both hiring and job posting measures of demand for early-career workers, we estimate joint-exposure difference-in-differences designs to disentangle these effects. Our findings point strongly towards WFH exposure as a better predictor of the decline in relative early-career hiring. We also find that actual WFH adoption (as opposed to occupational-exposure) also strongly predicts a fall in the share of new hires going to juniors and job postings requiring no more than three years experience.

Understanding the cause of the fall in relative demand for junior workers is important, as it has both scarring effects on the current cohort, and reduces long-run accumulation of human capital, leading to a slowdown in productivity. If GenAI is to blame, more drastic interventions, such as wage subsidies or differential tax treatment, may be needed to offset GenAI displacement effects. If, as our results indicate, WFH is reducing the incentive to hire junior talent, then more micro-level adjustments may be required to help firms adapt their organizational practices, so as to enjoy the benefits of WFH arrangements while simultaneously managing the development of early-career talent. Future work will require firm-level GenAI adoption measures with independent variation from WFH, as well as a longer post-period to track how these effects evolve over time.

References

ACEMOGLU, D. and J.-S. PISCHKE (1998). "Why Do Firms Train? Theory and Evidence". *The Quarterly Journal of Economics*, 113(1), 79–119.

ADÃO, R., M. KOLESÁR, and E. MORALES (Nov. 1, 2019). "Shift-Share Designs: Theory and Inference". *The Quarterly Journal of Economics*, 134(4), 1949–2010.

AKAN, M., J. M. BARRERO, N. BLOOM, T. BOWEN, S. BUCKMAN, S. DAVIS, and H. KIM (Mar. 2025). *The New Geography of Labor Markets*. NBER Working Paper w33582. Cambridge, MA: National Bureau of Economic Research.

AKSOY, C. G., N. BLOOM, S. DAVIS, V. MARINO, and C. ÖZGÜZEL (May 2026). *In-Person Contact Improves Remote-Work Performance*. Unreleased Working Paper.

AKSOY, C. G., N. BLOOM, S. J. DAVIS, V. MARINO, and C. ÖZGÜZEL (May 2025). *Remote Work, Employee Mix, and Performance*. NBER Working Paper w33851. Cambridge, MA: National Bureau of Economic Research.

ALMEIDA, L. (June 25, 2025). "UK Graduates Facing Worst Job Market since 2018 amid Rise of AI, Says Indeed". *The Guardian. Money*.

ALTHOFF, L. and H. REICHARDT (2026). *Task-Specific Technical Change and Comparative Advantage*. CESifo Working Paper 12403. Munich: CESifo.

ALTONJI, J. G., T. E. ELDER, and C. R. TABER (Apr. 1, 2008). "Using Selection on Observed Variables to Assess Bias from Unobservables When Evaluating Swan-Ganz Catheterization". *American Economic Review*, 98(2), 345–350.

APPEL, R., P. MCCRORY, A. TAMKIN, M. MCCAIN, T. NEYLON, and M. STERN (2025). *Anthropic Economic Index Report: Uneven Geographic and Enterprise AI Adoption*. Version 1. URL: https://arxiv.org/abs/2511.15080. Pre-published.

ARROW, K. J. (1962). "The Economic Implications of Learning by Doing". *The Review of Economic Studies*, 29(3), 155–173.

ATKIN, D., A. SCHOAR, and S. SHINDE (July 2023). *Working from Home, Worker Sorting and Development*. NBER Working Paper w31515. Cambridge, MA: National Bureau of Economic Research.

AZAR, J., M. GINE, and J. SANZ-ESPÍN (Dec. 1, 2025). *AI Is Already Eroding Wages: Quasi-Experimental Evidence From Occupational Exposure*. SSRN working paper 5842084. Rochester, NY: SSRN.

BARRERO, J. M., N. BLOOM, and S. J. DAVIS (Apr. 2021). *Why Working from Home Will Stick*. NBER Working Paper w28731. Cambridge, MA: National Bureau of Economic Research.

BARRON, J. M., D. A. BLACK, and M. A. LOEWENSTEIN (Jan. 1989). "Job Matching and On-the-Job Training". *Journal of Labor Economics*, 7(1), 1–19.

BATTISTON, D., J. BLANES I VIDAL, and T. KIRCHMAIER (Mar. 22, 2021). "Face-to-Face Communication in Organizations". *The Review of Economic Studies*, 88(2), 574–609.

BECKER, G. S. (1964). *Human Capital: A Theoretical and Empirical Analysis, with Special Reference to Education*. New York: Columbia University Press.

BEN-PORATH, Y. (1967). "The Production of Human Capital and the Life Cycle of Earnings". *Journal of Political Economy*, 75, (4, Part 1), 352–365.

BLATTER, M., S. MUEHLEMANN, and S. SCHENKER (Jan. 2012). "The Costs of Hiring Skilled Workers". *European Economic Review*, 56(1), 20–35.

BLOOM, N., G. B. DAHL, and D.-O. ROOTH (2026). "Work from Home and Disability Employment". *American Economic Review: Insights*, forthcoming.

BLOOM, N., J. LIANG, J. ROBERTS, and Z. J. YING (2015). "Does Working from Home Work? Evidence from a Chinese Experiment". *The Quarterly Journal of Economics*, 130(1), 165–218.

BORUSYAK, K., P. HULL, and X. JARAVEL (Jan. 10, 2022). "Quasi-Experimental Shift-Share Research Designs". *The Review of Economic Studies*, 89(1), 181–213.

BRYNJOLFSSON, E., B. CHANDAR, and R. CHEN (2025a). "Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence". *Stanford Digital Economy Lab Working Paper, November 2025*.

BRYNJOLFSSON, E., D. LI, and L. RAYMOND (Apr. 8, 2025b). "Generative AI at Work". *The Quarterly Journal of Economics*, 140(2), 889–942.

CALLAWAY, B., A. GOODMAN-BACON, and P. H. C. SANT'ANNA (2025). *Difference-in-Differences with a Continuous Treatment*. arXiv preprint 2107.02637v8. Version 8. arXiv.

CALVO, G. A. and S. WELLISZ (1978). "Supervision, Loss of Control, and the Optimum Size of the Firm". *Journal of Political Economy*, 86(5), 943–952.

CINELLI, C. and C. HAZLETT (Feb. 1, 2020). "Making Sense of Sensitivity: Extending Omitted Variable Bias". *Journal of the Royal Statistical Society Series B: Statistical Methodology*, 82(1), 39–67.

CULLEN, Z., B. PAKZAD-HURSON, and R. PEREZ-TRUGLIA (May 1, 2025). "Home Sweet Home: How Much Do Employees Value Remote Work?" *AEA Papers and Proceedings*, 115, 276–281.

DE CHAISEMARTIN, C., X. D'HAULTFŒUILLE, and G. VAZQUEZ-BARE (May 1, 2024). "Difference-in-Difference Estimators with Continuous Treatments and No Stayers". *AEA Papers and Proceedings*, 114, 610–613.

DEMIRER, M., J. HORTON, N. IMMORLICA, B. LUCIER, and P. SHAHIDI (Feb. 2026). *Chaining Tasks, Redefining Work: A Theory of AI Automation*. NBER Working Paper w34859. Cambridge, MA: National Bureau of Economic Research.

DINGEL, J. I. and B. NEIMAN (Sept. 1, 2020). "How Many Jobs Can Be Done at Home?" *Journal of Public Economics*, 189, 104235.

ELOUNDOU, T., S. MANNING, P. MISHKIN, and D. ROCK (June 21, 2024). "GPTs Are GPTs: Labor Market Impact Potential of LLMs". *Science*, 384(6702), 1306–1308.

EMANUEL, N., E. HARRINGTON, and A. PALLAIS (June 11, 2024). "Research: How Remote Work Impacts Women at Different Stages of Their Careers". *Harvard Business Review*.

— (May 12, 2026). "The Power of Proximity to Coworkers". *The Quarterly Journal of Economics*, qjag027.

FARBER, H. S. and R. GIBBONS (Nov. 1, 1996). "Learning and Wage Dynamics". *The Quarterly Journal of Economics*, 111(4), 1007–1047.

FELTEN, E., M. RAJ, and R. SEAMANS (2023). *How Will Language Modelers like ChatGPT Affect Occupations and Industries?* Version 2. URL: https://arxiv.org/abs/2303.01157. Pre-published.

GANS, J. and A. GOLDFARB (Jan. 2026). *O-Ring Automation*. NBER Working Paper w34639. Cambridge, MA: National Bureau of Economic Research.

GARICANO, L. (2000). "Hierarchies and the Organization of Knowledge in Production". *Journal of Political Economy*, 108(5), 874–904.

GARICANO, L., J. LI, and Y. WU (2026). "Weak Bundle, Strong Bundle: How AI Redraws Job Boundaries". *CEPR Discussion Paper*.

GARICANO, L. and E. ROSSI-HANSBERG (Nov. 2006). "Organization and Inequality in a Knowledge Economy". *Quarterly Journal of Economics*, 121(4), 1383–1435.

GIBBONS, R. and M. WALDMAN (1999). "A Theory of Wage and Promotion Dynamics inside Firms". *The Quarterly Journal of Economics*, 114(4), 1321–1358.

GILLEN, B., E. SNOWBERG, and L. YARIV (Aug. 2019). "Experimenting with Measurement Error: Techniques with Applications to the Caltech Cohort Study". *Journal of Political Economy*, 127(4), 1826–1863.

HANSEN, S., P. J. LAMBERT, N. BLOOM, S. J. DAVIS, R. SADUN, and B. TASKA (Mar. 2023). *Remote Work across Jobs, Companies, and Space*. NBER Working Paper w31007. Cambridge, MA: National Bureau of Economic Research.

HARRIS, M. and B. HOLMSTRÖM (July 1982). "A Theory of Wage Dynamics". *The Review of Economic Studies*, 49(3), 315–333.

HARTLEY, J., F. JOLEVSKI, V. MELO, and B. MOORE (2025). *The Labor Market Effects of Generative Artificial Intelligence*. SSRN working paper 5136877. SSRN.

HASHIMOTO, M. (1981). "Firm-Specific Human Capital as a Shared Investment". *American Economic Review*, 71(3), 475–482.

HO, D. E., K. IMAI, G. KING, and E. A. STUART (2007). "Matching as Nonparametric Preprocessing for Reducing Model Dependence in Parametric Causal Inference". *Political Analysis*, 15(3), 199–236.

HOLMSTRÖM, B. (1982). "Moral Hazard in Teams". *The Bell Journal of Economics*, 13(2), 324–340.

HOLMSTRÖM, B. (Jan. 1999). "Managerial Incentive Problems: A Dynamic Perspective". *The Review of Economic Studies*, 66(1), 169–182.

HOSSEINI MAASOUM, S. M. and G. LICHTINGER (Aug. 31, 2025). *Generative AI as Seniority-Biased Technological Change: Evidence from U.S. Résumé and Job Posting Data*. SSRN Scholarly Paper 5425555. Rochester, NY: SSRN.

HUMLUM, A. and E. VESTERGAARD (May 2025). *Still Waters, Rapid Currents: Early Labor Market Transformation under Generative AI*. NBER Working Paper w33777. Cambridge, MA: National Bureau of Economic Research.

IACUS, S. M., G. KING, and G. PORRO (2012). "Causal Inference without Balance Checking: Coarsened Exact Matching". *Political Analysis*, 20(1), 1–24.

JAROSCH, G., E. OBERFIELD, and E. ROSSI-HANSBERG (2021). "Learning From Coworkers". *Econometrica*, 89(2), 647–676.

KHARAZIAN, A. (2026). *Ramp AI Index*. URL: https://ramp.com/data/ai-index.

KWAN, A., B. MATTHIES, R. TOWNSEND, and T. XU (May 2025). *Entrepreneurial Spawning from Remote Work*. NBER Working Paper w33774. Cambridge, MA: National Bureau of Economic Research.

LAMBERT, P. J. and C. LARKIN (2024). *Has Work from Home Shifted the US Electoral Map?* CEPOP67. Centre for Economic Performance, LSE.

LAZEAR, E. P. and P. OYER (Dec. 31, 2013). "Personnel Economics". In: *The Handbook of Organizational Economics*. Ed. by R. GIBBONS and J. ROBERTS. Princeton: Princeton University Press, pp. 479–519.

LAZEAR, E. P., K. L. SHAW, and C. T. STANTON (Oct. 2015). "The Value of Bosses". *Journal of Labor Economics*, 33(4), 823–861.

MANNING, A. (Jan. 1, 2006). "A Generalised Model of Monopsony". *The Economic Journal*, 116(508), 84–100.

MAS, A. and A. PALLAIS (Dec. 1, 2017). "Valuing Alternative Work Arrangements". *American Economic Review*, 107(12), 3722–3759.

MCELHERAN, K., J. F. LI, E. BRYNJOLFSSON, Z. KROFF, E. DINLERSOZ, L. FOSTER, and N. ZOLAS (Mar. 2024). "AI Adoption in America: Who, What, and Where". *Journal of Economics & Management Strategy*, 33(2), 375–415.

MINCER, J. (1974). *Schooling, Experience, and Earnings*. National Bureau of Economic Research, Inc.

NOY, S. and W. ZHANG (July 14, 2023). "Experimental Evidence on the Productivity Effects of Generative Artificial Intelligence". *Science*, 381(6654), 187–192.

OREOPOULOS, P., T. VON WACHTER, and A. HEISZ (Jan. 1, 2012). "The Short- and Long-Term Career Effects of Graduating in a Recession". *American Economic Journal: Applied Economics*, 4(1), 1–29.

OSTER, E. (Apr. 3, 2019). "Unobservable Selection and Coefficient Stability: Theory and Evidence". *Journal of Business & Economic Statistics*, 37(2), 187–204.

PALLAIS, A. (Nov. 1, 2014). "Inefficient Hiring in Entry-Level Labor Markets". *American Economic Review*, 104(11), 3565–3599.

PAPKE, L. E. and J. M. WOOLDRIDGE (1996). "Econometric Methods for Fractional Response Variables with an Application to 401(k) Plan Participation Rates". *Journal of Applied Econometrics*, 11(6), 619–632.

PFANN, G. A. and B. VERSPAGEN (1989). "The Structure of Adjustment Costs for Labour in the Dutch Manufacturing Sector". *Economics Letters*, 29(4), 365–371.

PRENDERGAST, C. (Mar. 1, 1999). "The Provision of Incentives in Firms". *Journal of Economic Literature*, 37(1), 7–63.

QIAN, Y. (July 1, 1994). "Incentives and Loss of Control in an Optimal Hierarchy". *The Review of Economic Studies*, 61(3), 527–544.

RAMAN, A. (May 19, 2025). "LinkedIn Executive: A.I. Is Coming for Entry-Level Jobs". *The New York Times. Opinion.*

ROBINSON, P. M. (1988). "Root-N-consistent Semiparametric Regression". *Econometrica*, 56(4), 931–954.

ROSEN, S. (Aut. 1982). "Authority, Control, and the Distribution of Earnings". *The Bell Journal of Economics*, 13(2), 311–323.

SARGENT, T. J. (Dec. 1978). "Estimation of Dynamic Labor Demand Schedules under Rational Expectations". *Journal of Political Economy*, 86(6), 1009–1044.

SATTINGER, M. (May 1975). "Comparative Advantage and the Distributions of Earnings and Abilities". *Econometrica*, 43(3), 455–468.

SCHWANDT, H. and T. VON WACHTER (2019). "Unlucky Cohorts: Estimating the Long-Term Effects of Entering the Labor Market in a Recession in Large Cross-Sectional Data Sets". *Journal of Labor Economics*, 37(S1), S161–S198.

SEN, R. (Mar. 2024). *Supervision at Work: Evidence from a Field Experiment*. Working Paper UGA-22057. London: International Growth Centre.

SHRIVASTAVA, A. (Sept. 25, 2025). *September 2025 Labor Market Update: The Squeeze on New Entrants Mirrors a Marketwide Decline*. Indeed Hiring Lab. URL: https://www.hiringlab.org/2025/09/25/september-labor-market-squeeze-on-new-entrants/.

STANTON, C. T. and C. THOMAS (Apr. 2016). "Landing the First Job: The Value of Intermediaries in Online Hiring". *The Review of Economic Studies*, 83(2), 810–854.

TEESELINK, B. K. (Dec. 2025). *Generative AI and Labor Market Outcomes: Evidence from the United Kingdom*. SSRN working paper 5516798. SSRN.

THOMPSON, D. (Apr. 30, 2025). *Something Alarming Is Happening to the Job Market*. The Atlantic. URL: https://www.theatlantic.com/economy/archive/2025/04/job-market-youth/682641/.

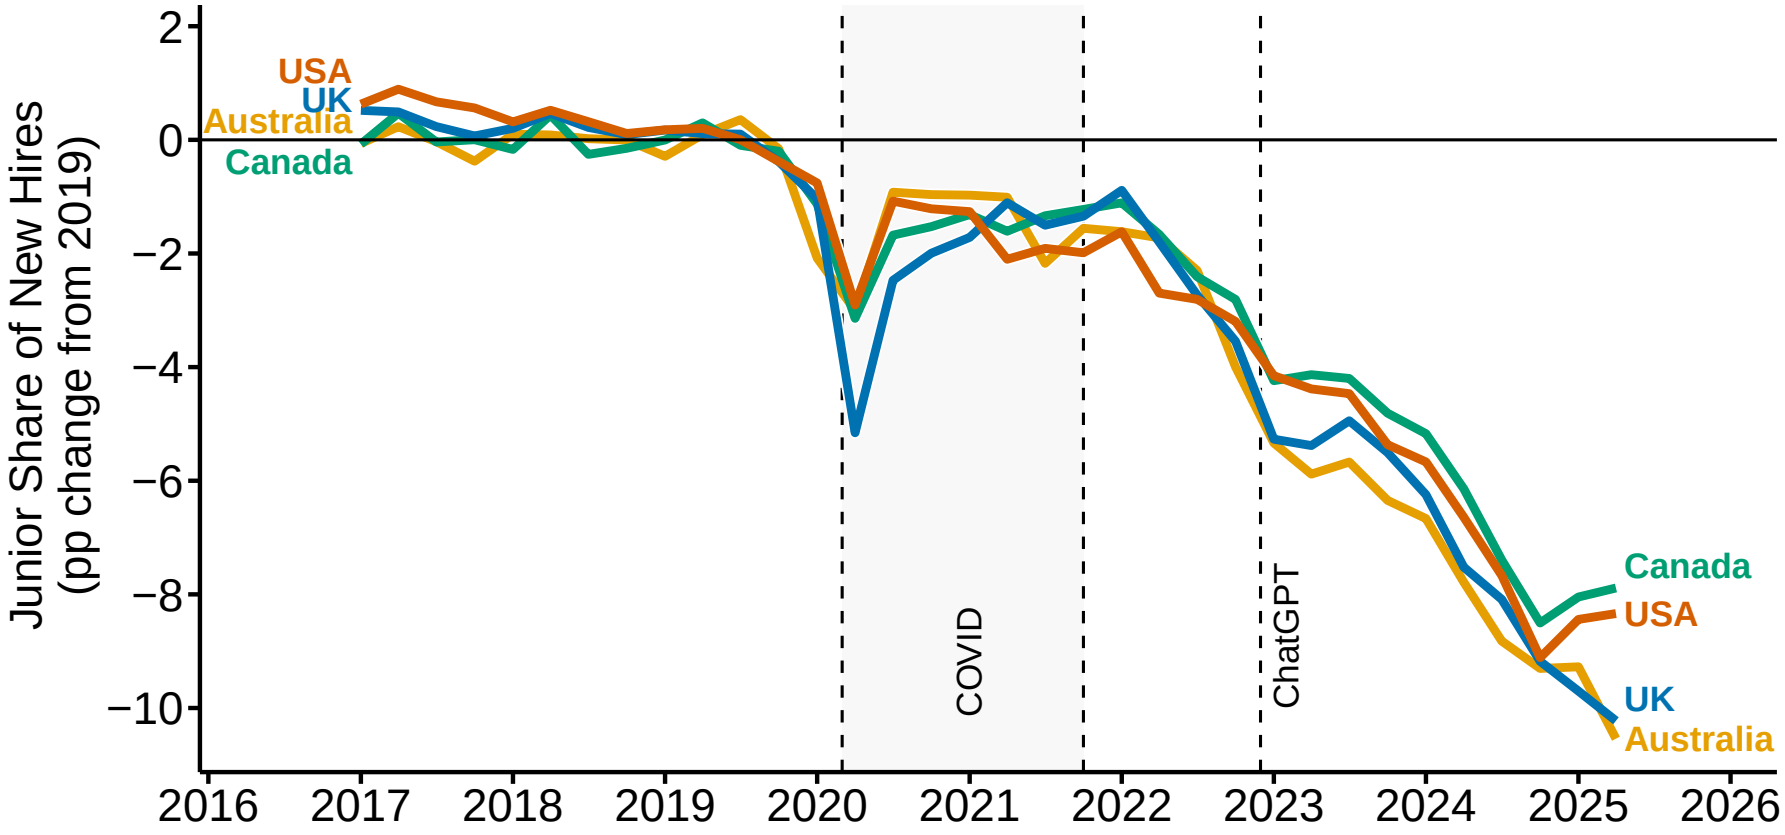
TOPEL, R. H. and M. P. WARD (1992). "Job Mobility and the Careers of Young Men". *The Quarterly Journal of Economics*, 107(2), 439–479.

WALDMAN, M. (Jan. 1984). "Worker Allocation, Hierarchies and the Wage Distribution". *The Review of Economic Studies*, 51(1), 95–109.

YANG, L., D. HOLTZ, S. JAFFE, S. SURI, S. SINHA, J. WESTON, C. JOYCE, N. SHAH, K. SHERMAN, B. HECHT, and J. TEEVAN (Jan. 2022). "The Effects of Remote Work on Collaboration among Information Workers". *Nature Human Behaviour*, 6(1), 43–54.

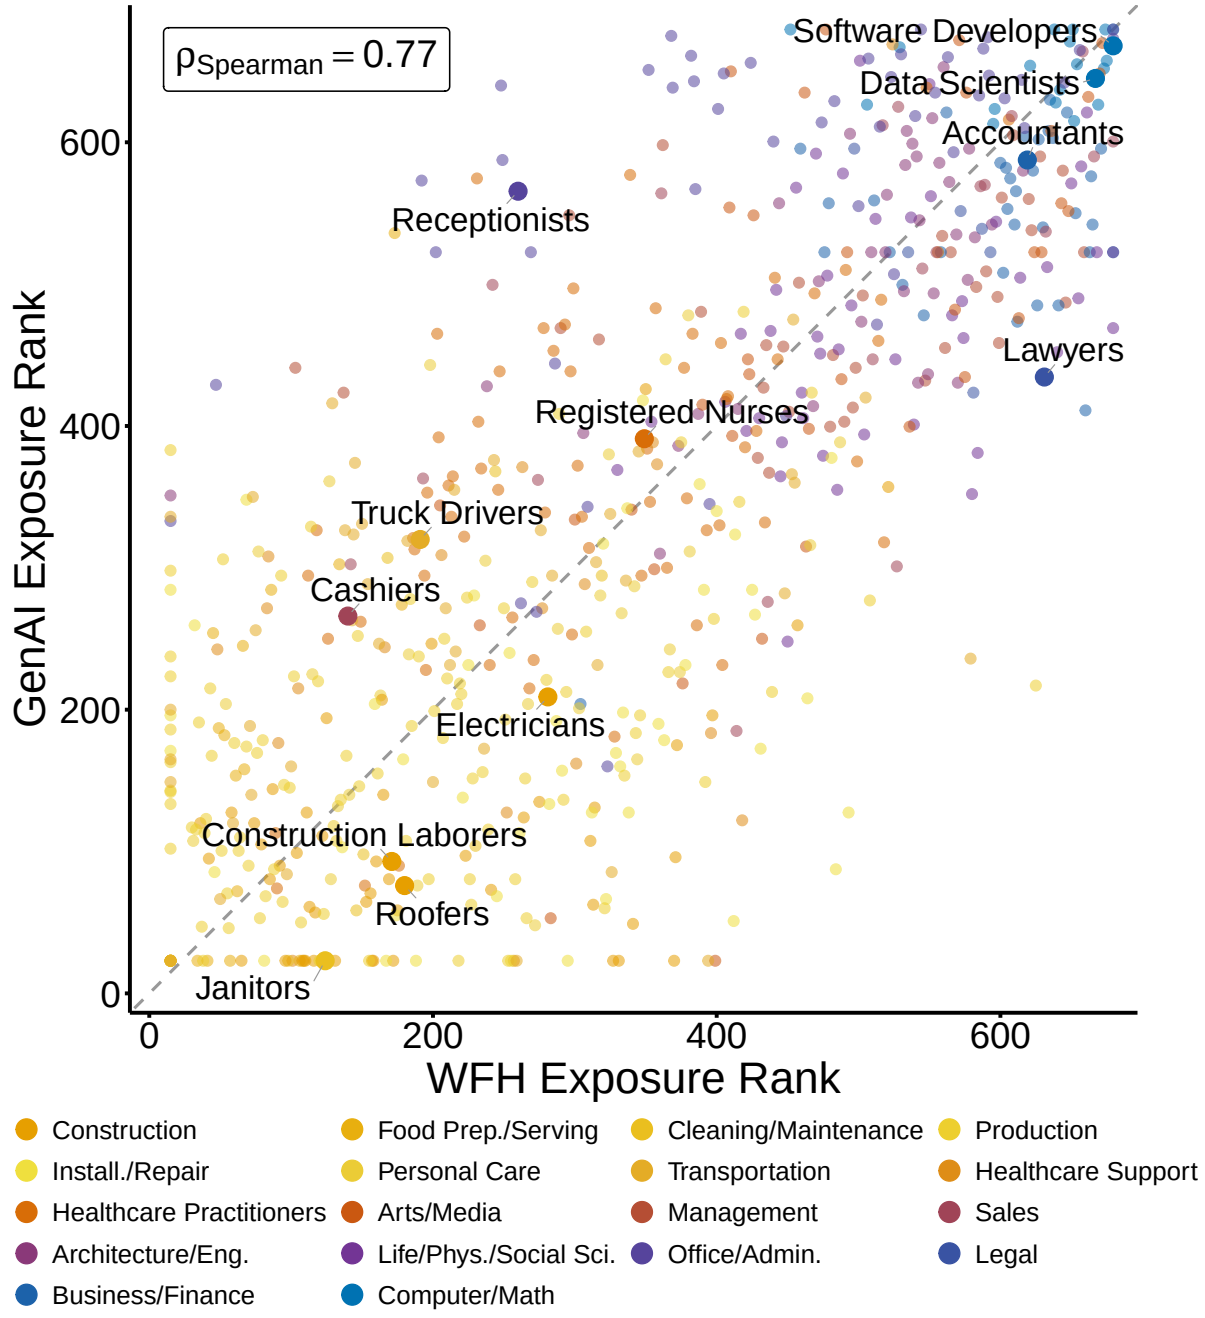
YOTZOV, I., J. M. BARRERO, N. BLOOM, P. BUNN, S. DAVIS, K. FOSTER, A. JALCA, B. MEYER, P. MIZEN, M. NAVARRETE, P. SMIETANKA, G. THWAITES, and B. Z. WANG (Feb. 2026). *Firm Data on AI*. NBER Working Paper w34836. Cambridge, MA: National Bureau of Economic Research.

Figure 1: Junior-share of Hiring Has Declined Sharply Across Four Anglophone Economies, 2017–2025



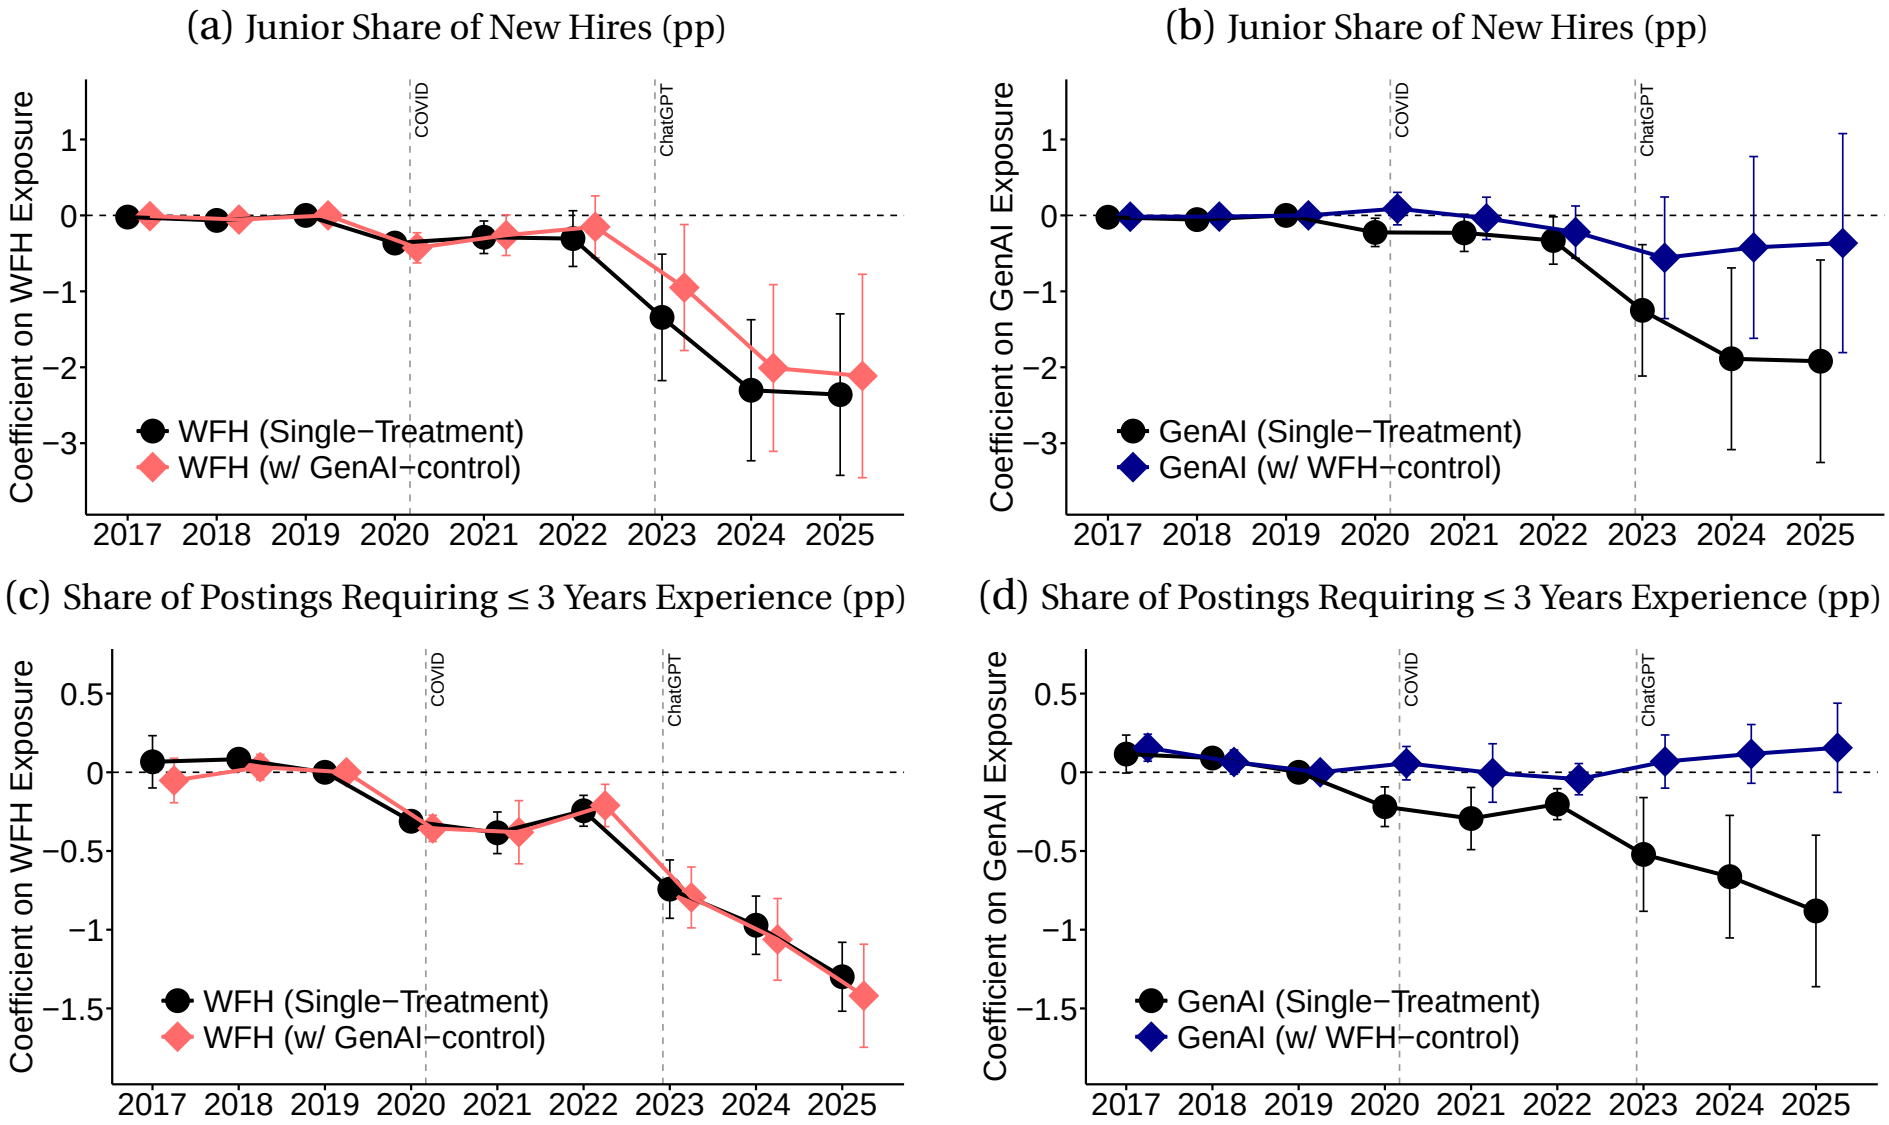
Note: Figure depicts the "junior-share" of new hires for the US, UK, Canada, and Australia, expressed as the percentage-point change from each country's 2019 mean. Data comes from Revelio Labs (see Section 4 for detailed description). "Junior" is defined using the bottom-two groups of a six-level scale which accounts for the job description and each worker's past experience at the beginning of their tenure in a new position. All series are constructed as a weighted arithmetic mean of 3-digit-occupation junior shares, using the US 2019 3-digit O*NET share of all hires as the (fixed) weight, to net out temporal and cross-country compositional differences. The 2019 (reweighted) junior share of new hires is 66.9% in the US, 64.9% in the UK, 69.3% in Canada, and 65.8% in Australia. Series are at quarterly frequency and seasonally adjusted using the U.S. Census Bureau's X-13ARIMA-SEATS procedure applied separately to each country's series. Dashed vertical lines mark the COVID-19 pandemic and the public release of ChatGPT.

Figure 2: WFH and GenAI Exposure Are Highly Correlated Across 6-Digit O*NET-SOC Occupations



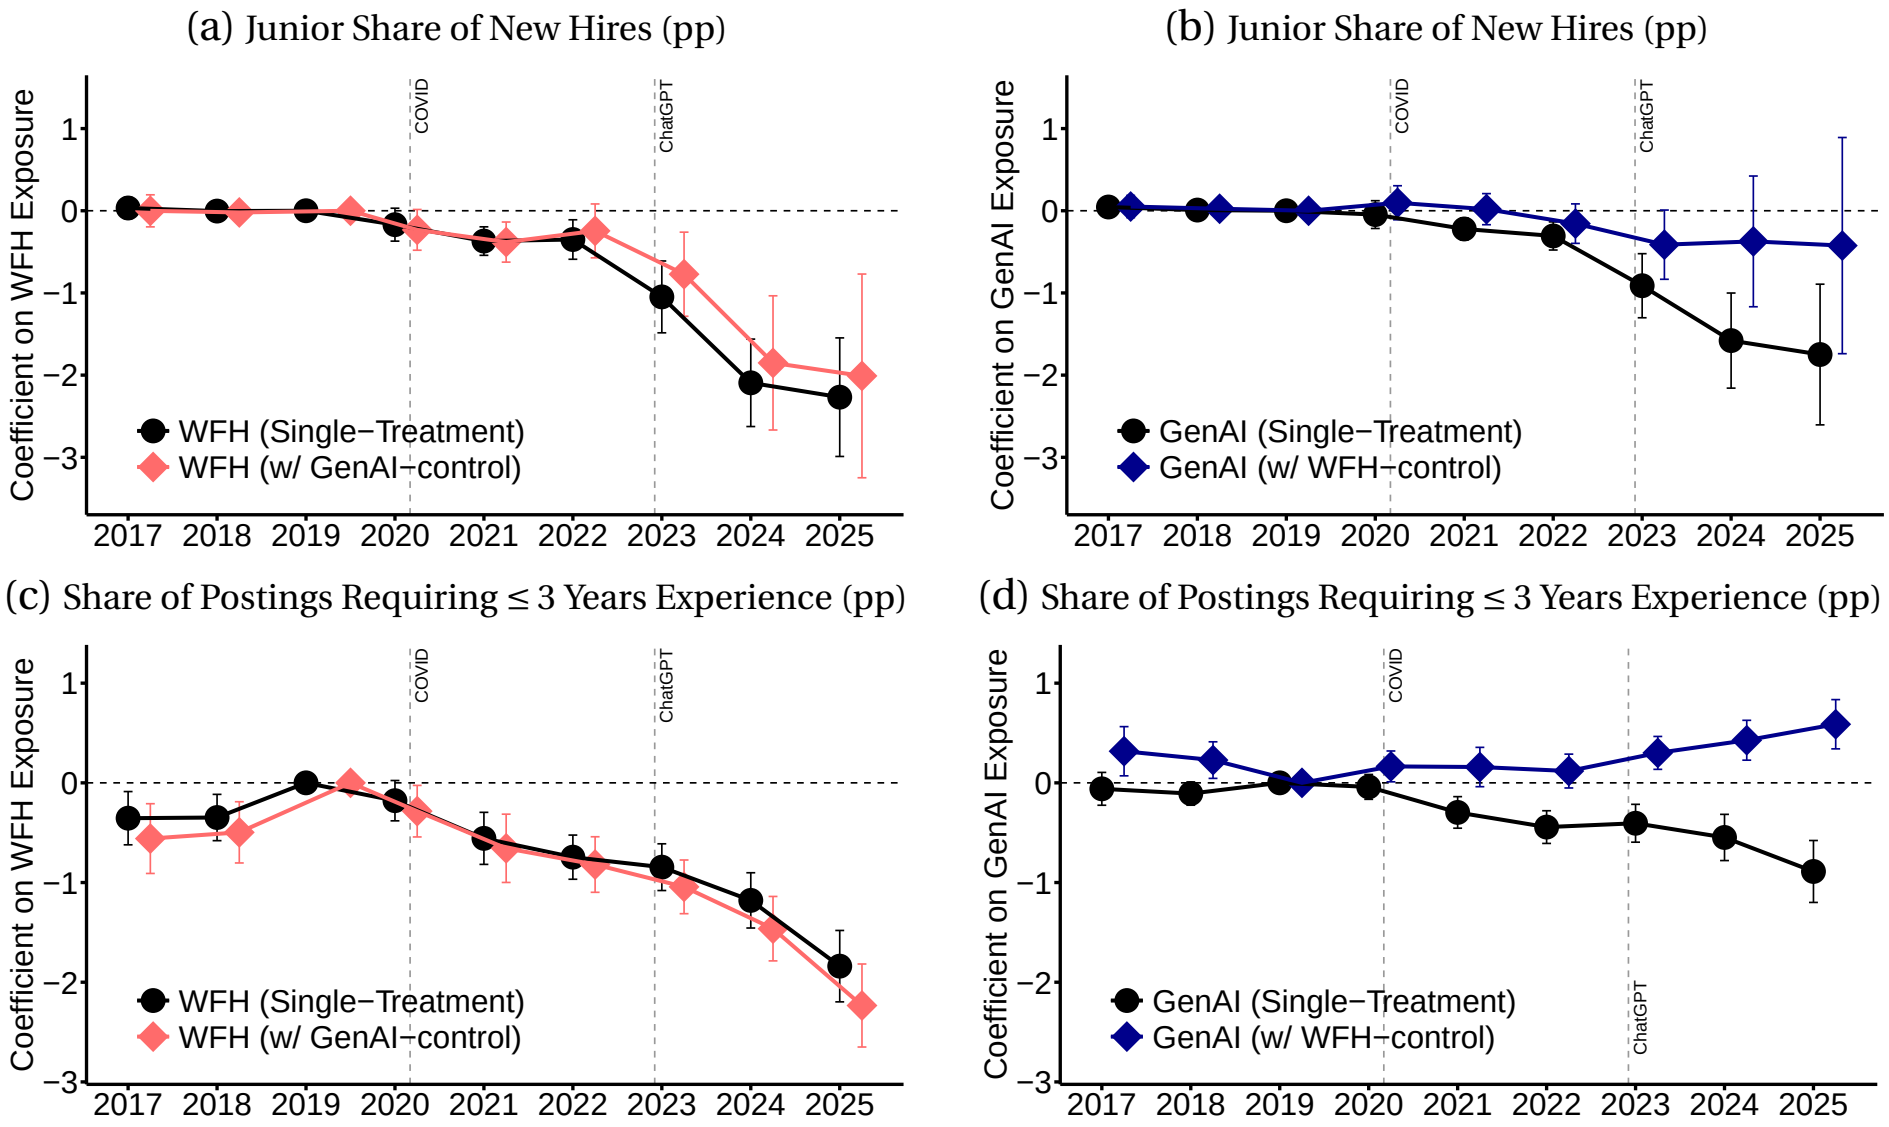
Note: This figure depicts 683 6-digit O*NET-SOC occupations showing their WFH exposure from Hansen et al. (2023) and their GenAI exposure from Eloundou et al. (2024) (see Section 4 for a discussion of both measures). The figure plots cross-sectional ranks of each occupation on the two indices; the dashed line is the 45-degree reference. The Spearman rank correlation is $\hat{\rho}_S = 0.77$, and the OLS regression of standardized GenAI on standardized WFH exposure gives $\hat{\beta} = 0.71$ (s.e. 0.03, $R^2 = 0.50$).

Figure 3: Measured WFH Exposure Remains Predictive After Controlling for Measured GenAI Exposure—the Reverse Does Not Hold (Firm×Broad-Occupation Units)



Note: Each panel plots year-by-year coefficients from Equation 8, from two regressions of the junior hiring outcome on (i) the treatment exposure alone (black) and (ii) the joint specification (red/blue). The underlying panel is firm×2-digit O*NET×country×year (2017–2025), with firm×occupation and country×year fixed effects, weighted by unit totals. Because exposure is a shift-share over 6-digit-occupation shocks, we report shock-level standard errors following Borusyak et al. (2022). Bars are 95% confidence intervals; 2019 is the reference year. Panels (a) and (b) report the junior share of new hires (pp); panels (c) and (d) the share of postings requiring ≤ 3 years experience (pp).

Figure 4: The WFH–GenAI Asymmetry Replicates at the Occupation×State Level



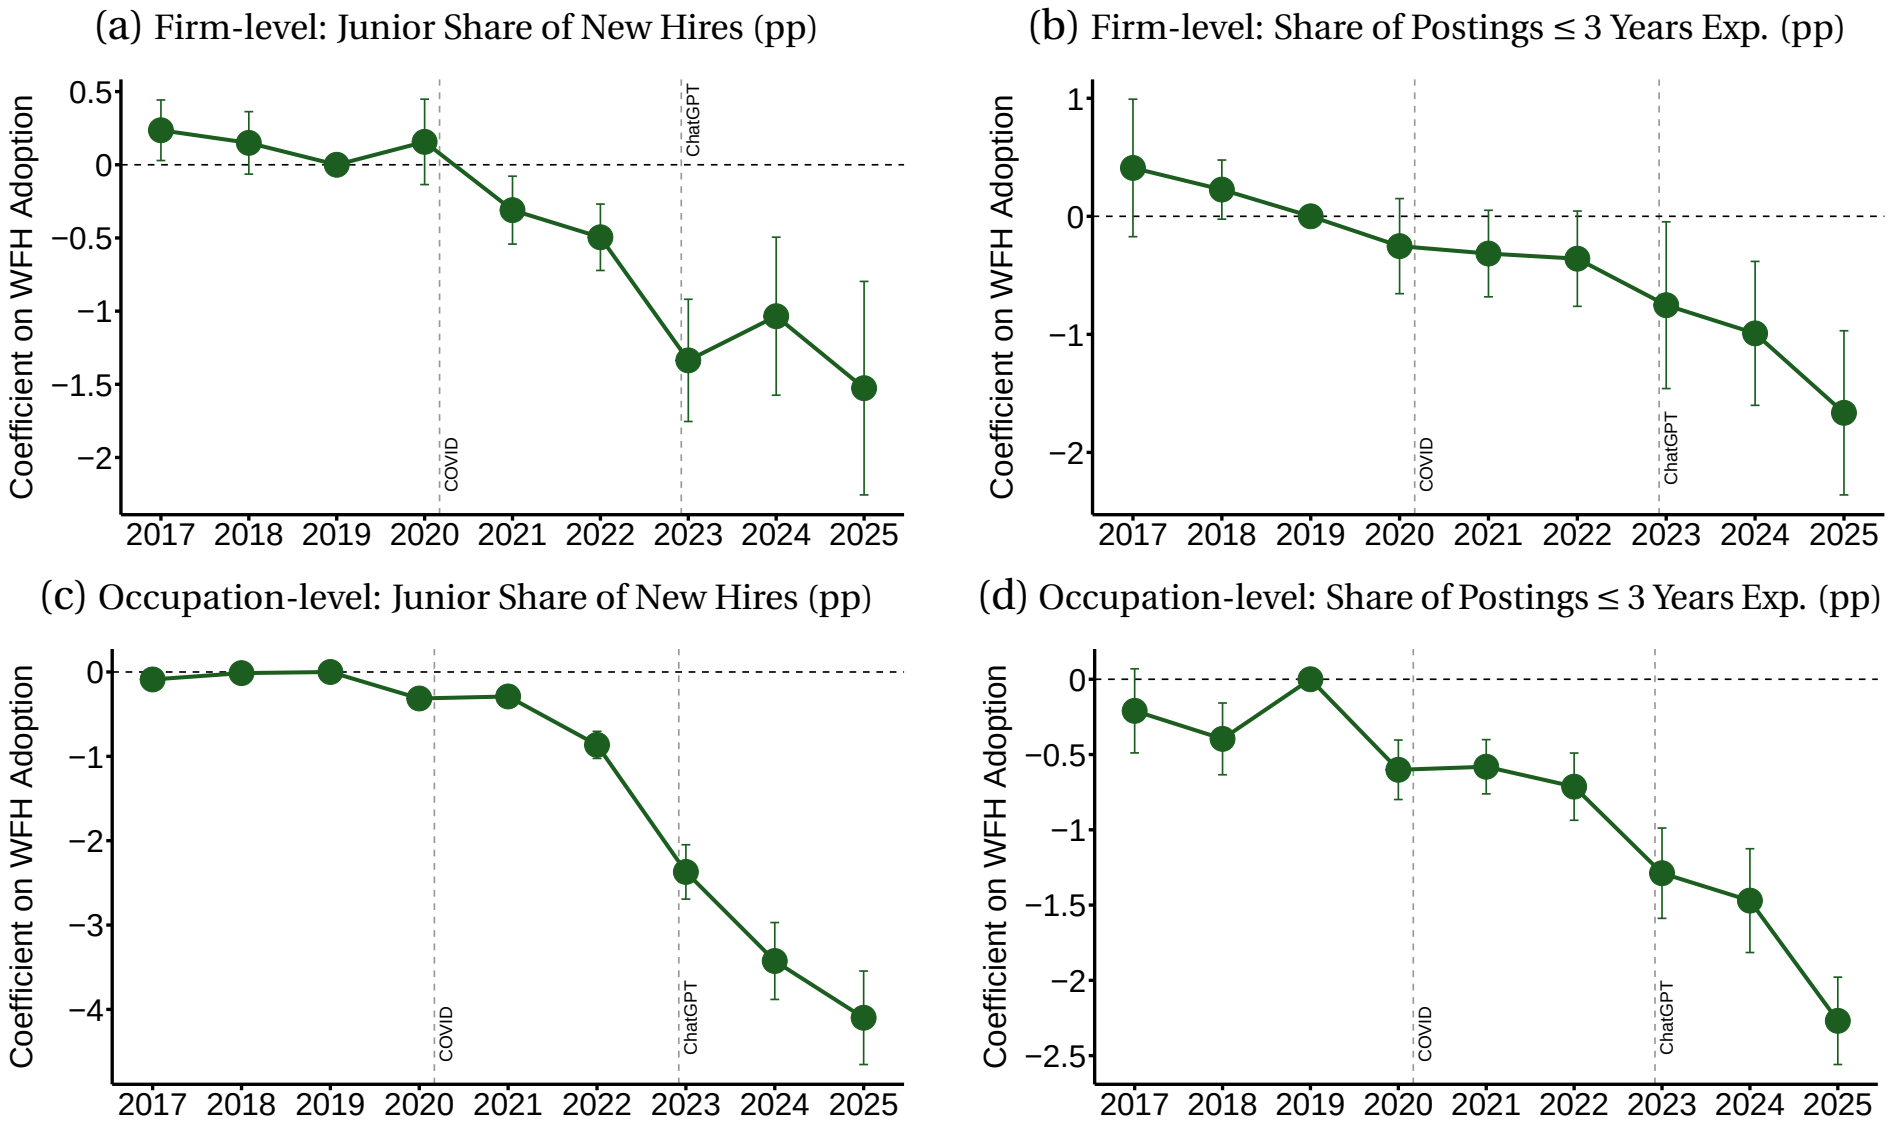
Note: Companion specification to Figure 3 estimating Equation 8 at the occupation×state×month level rather than firm×occupation×year. Each panel plots year-by-year coefficients from two separate regressions of the junior hiring outcome on (i) the treatment exposure alone (black), and (ii) the treatment exposure jointly with the co-treatment (red/blue). The unit of observation is 6-digit O*NET-SOC occupation×state×calendar month, with state×occupation and state×month fixed effects (where state denotes country×state), weighted by unit totals. Standard errors are two-way clustered by occupation and state, with 95% confidence intervals shown and 2019 as the reference year. Panels (a) and (b) report the junior share of new hires (pp) and panels (c) and (d) the share of postings requiring ≤ 3 years experience (pp).

Table 1: Measured WFH exposure remains predictive of post-2022 declines in junior hiring and recruitment under joint specification

	Junior Worker Share of New Hires (Pp.)				Share of Job Postings Requiring ≤ 3 Years Experience (Pp.)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: State × Occupation × Month Cells								
Post × WFH (1sd)	-1.69***		-1.42***	-1.63***	-1.06***		-1.24***	-1.35***
	(0.24)		(0.36)	(0.37)	(0.14)		(0.18)	(0.27)
Post × GenAI (1sd)		-1.34***	-0.41	0.04		-0.56***	0.27**	0.35*
		(0.27)	(0.36)	(0.42)		(0.12)	(0.11)	(0.19)
Post × WFH × GenAI (1sd)				0.66***				0.15
				(0.23)				(0.22)
Partial *F*-Statistic	48.7	25.5	15.8/1.3	19.0/0.0/8.2	56.2	21.5	48.3/5.7	25.0/3.5/0.5
Variance Inflation Factor	–	–	1.8	1.9/2.2/1.3	–	–	1.8	2.9/2.5/1.7
Condition number	–	–	2.2	2.6	–	–	2.2	3.1
Observations	47,003 units / 2,312,186 cells / 120,659,812 hires				207,428 units / 6,577,807 cells / 217,318,081 postings			
Panel B: Firm × 2-Digit Occupation × Year Cells								
Post × WFH (1sd)	-1.88***		-1.57***	-1.57***	-1.10***		-1.15***	-1.10***
	(0.42)		(0.45)	(0.42)	(0.11)		(0.18)	(0.16)
Post × GenAI (1sd)		-1.59***	-0.45	-0.45		-0.80***	0.06	0.04
		(0.51)	(0.49)	(0.55)		(0.21)	(0.12)	(0.11)
Post × WFH × GenAI (1sd)				-0.01				-0.05
				(0.28)				(0.06)
Partial *F*-Statistic	20.6	9.9	12.1/0.8	13.6/0.7/0.0	93.8	14.5	40.7/0.2	47.9/0.1/0.7
Variance Inflation Factor	–	–	2.0	2.1/2.2/1.1	–	–	2.3	2.8/2.5/1.2
Condition number	–	–	2.5	2.6	–	–	2.7	3.1
Observations	603,650 units / 3,621,900 cells / 31,706,183 hires				344,211 units / 2,065,266 cells / 58,940,327 postings			

Note: Effect of occupational GenAI and WFH exposure on early-career hiring across the US, UK, Australia, and Canada, estimating (9). Columns (1)–(4) use the junior share of new hires, and columns (5)–(8) the share of postings requiring ≤ 3 years experience. Two-period DiD specifications follow Section 5, with pre-period 2017–2019 and post-period 2023–2025. *Panel A* uses State×6-digit-Occupation×Month cells, with State×Occupation and State×Month fixed effects, two-way clustered by occupation and state. *Panel B* uses Firm×2-digit-Occupation×Year cells with shift-share robust inference following Borusyak et al. (2022), and exposure built from each unit's 2017–2019 occupation mix. VIF, CN, and partial *F* are reported in the same order as the coefficients (see Appendix F.2 for discussion of these multi-collinearity diagnostics). Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Figure 5: Actual WFH Adoption Predicts Lower Early-Career Hiring and Recruitment, Even After Matching on WFH and GenAI Exposure



Note: Year-by-year coefficients on the interaction of actual WFH adoption with year, relative to 2019, from the event-study analogues of Equation 10 (panels (a) and (b)) and Equation 11 (panels (c) and (d)). Panels (a) and (b) use a binary firm-level treatment (1 if the firm posted any Hansen et al., 2023-classified WFH job in 2021–22), with firm×2-digit-occupation and country×year fixed effects, and SEs two-way clustered by firm and country×2-digit-occupation. Panels (c) and (d) use the standardized 2021–22 WFH share of postings, with country×state×6-digit-occupation unit and country×year fixed effects, and SEs clustered by unit, on the same activity-window sample. All specifications include year-interacted log-size controls. Bars are 95% confidence intervals, and dashed lines mark COVID onset and the ChatGPT release.

Table 2: Actual WFH adoption predicts lower intensity of early-career hiring and recruitment, even after matching on exposure

	Junior Worker Share of New Hires (Pp.)				Share of Job Postings Requiring ≤ 3 Years Experience (Pp.)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Firm × 2-Digit Occupation × Year (binary actual-WFH treatment)								
Post × 1[Any firm WFH 2021–22]	-1.20***	-0.94***	-1.04***	-0.94***	-1.02***	-0.69***	-0.81***	-0.70***
	(0.20)	(0.15)	(0.16)	(0.16)	(0.26)	(0.22)	(0.23)	(0.22)
Post × Occ. WFH (1sd)		-2.59***		-2.81***		-0.98***		-1.30***
		(0.29)		(0.64)		(0.14)		(0.21)
Post × Occ. GenAI (1sd)			-1.94***	0.27			-0.67***	0.37**
			(0.49)	(0.60)			(0.19)	(0.16)
Partial F-Statistic	36.6	37.3/77.7	42.4/15.9	36.6/19.5/0.2	14.9	9.7/48.7	12.8/12.1	10.2/38.5/5.4
Variance Inflation Factor	–	1.0	1.0	1.0/2.8/2.8	–	1.0	1.0	1.0/3.1/3.1
Condition number	–	1.0	1.0	3.0	–	1.0	1.0	3.2
Observations	222,452 units / 1,955,646 cells / 25,037,345 hires				326,000 units / 2,854,065 cells / 90,262,954 postings			
Panel B: Country × State × 6-Digit Occupation × Year (continuous actual-WFH treatment)								
Post × Realised WFH share (1sd)	-2.91***	-0.06	-3.10***	-0.16	-1.23***	-0.64***	-1.24***	-0.64***
	(0.20)	(0.27)	(0.23)	(0.27)	(0.14)	(0.24)	(0.19)	(0.24)
Post × Occ. WFH (1sd)		-2.58***		-2.73***		-0.53**		-0.55**
		(0.25)		(0.27)		(0.22)		(0.25)
Post × Occ. GenAI (1sd)			0.54***	0.79***			0.03	0.06
			(0.18)	(0.19)			(0.16)	(0.17)
Partial F-Statistic	219.8	0.0/104.6	174.4/8.7	0.4/103.6/17.7	74.3	7.0/5.6	42.7/0.0	7.1/4.8/0.1
Variance Inflation Factor	–	6.6	1.3	6.7/6.8/1.3	–	12.4	1.2	12.4/12.6/1.2
Condition number	–	4.9	1.7	5.4	–	6.9	1.6	7.4
Observations	44,277 units / 518,187 cells / 225,724,747 hires				44,240 units / 377,760 cells / 341,664,290 postings			

Note: Effect of actual WFH adoption on early-career hiring in the US, UK, Australia, and Canada, estimating (10) in Panel A and (11) in Panel B. Cols. (1)–(4) use the junior share of new hires; cols. (5)–(8) the share of postings requiring ≤ 3 years experience. Specifications use Coarsened Exact Matching (CEM) on a 5×5 grid of WFH/GenAI exposure quintiles (Section 5); treatments are defined in Section 4. *Panel A*: Firm×2-digit-Occupation×Year cells, binary treatment, firm×occupation and country×year FEs, a post-interacted pre-period firm log-size control, and SEs two-way clustered by firm and country×2-digit-occupation. *Panel B*: Country×State×6-digit-Occupation×Year cells, continuous treatment, cell and country×year FEs, a post-interacted pre-period unit log-size control, and SEs clustered by cell. The sample is units with > 1 posting/hire in each of 2017–19, 2020–21, and 2023–25 (Panel A requires a postings–hires crosswalk match). Col. (1) reports the actual-WFH coefficient alone; cols. (2)/(3) add the Hansen WFH or Eloundou GenAI index; col. (4) adds both. VIF, CN, and partial F follow the coefficient order (see Appendix F.2). Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

The Broken Ladder:
AI, Remote Work, and Early-Career Hiring

Appendix: For Online Publication Only

A Analytical Derivations for the Partial-Equilibrium Model

This appendix proves the comparative statics (5) and (6) from Section 3. The comparative statics are fixed-worker-stock comparative statics. For every parameter θ considered below, the current employment stocks used to evaluate production are held fixed:

$$\frac{\partial \ell^S}{\partial \theta} = 0, \qquad \frac{\partial \ell^J}{\partial \theta} = 0. \tag{A.1}$$

Thus the derivatives below are partial derivatives of continuation values and hiring shares at a given team composition. They are not total derivatives after allowing the current worker stocks to adjust, and they do not impose additional stock-flow steady-state restrictions when the parameter changes. Only the effective inputs

$$s \equiv \ell^S - \phi \ell^J, \qquad j \equiv \eta \ell^J \tag{A.2}$$

move mechanically with parameters. In particular,

$$\frac{\partial s}{\partial \phi} = -\ell^J, \qquad \frac{\partial j}{\partial \phi} = 0, \qquad \frac{\partial s}{\partial \eta} = 0, \qquad \frac{\partial j}{\partial \eta} = \ell^J, \qquad \frac{\partial s}{\partial \lambda} = \frac{\partial j}{\partial \lambda} = 0. \tag{A.3}$$

Throughout, $s > 0$, $\ell^J > 0$, and output is

$$Y = AF(s, j). \tag{A.4}$$

The production function F is twice continuously differentiable, strictly increasing, concave, homogeneous of degree one, strictly concave in the factor ratio, satisfies Inada conditions, and has finite positive local elasticity of substitution. All production derivatives are evaluated at the fixed-stock point (s, j).

A.1 Marginal products

Operating profit before hiring costs is

$$\pi = AF(s, j) - w_S \ell^S - w_J \ell^J. \tag{A.5}$$

The current marginal products entering the value equations are the fixed-current-stock marginal products of adding a worker of each type. Since $\partial s/\partial \ell^S = 1$, the marginal product of a senior is

$$\mathrm{MP}_S = AF_s(s, j). \tag{A.6}$$

Since $\partial s/\partial \ell^J = -\phi$ and $\partial j/\partial \ell^J = \eta$, the marginal product of a junior is

$$\mathrm{MP}_J = A\left[\eta F_j(s, j) - \phi F_s(s, j)\right]. \tag{A.7}$$

Define

$$M \equiv \mathrm{MP}_S = AF_s(s, j), \qquad \psi \equiv \frac{\mathrm{MP}_J}{\mathrm{MP}_S} = \eta \frac{F_j(s, j)}{F_s(s, j)} - \phi. \tag{A.8}$$

Then $\mathrm{MP}_J = \psi M$. Because $F_s > 0$, $M > 0$. Under the maintained assumption $\mathrm{MP}_J > 0$, we also have $\psi > 0$.

A.2 Junior hiring share and value ratio

Let

$$B \equiv r + \delta_S, \qquad D \equiv r + \lambda + \delta_J. \tag{A.9}$$

Continuation values. We normalise the value of a separated worker to zero. A marginal senior produces current flow surplus $M - w_S$ and separates at rate δ_S. The senior Bellman equation is therefore

$$rV_S = M - w_S - \delta_S V_S. \tag{A.10}$$

Solving gives

$$V_S = \frac{M - w_S}{r + \delta_S} = \frac{M - w_S}{B}. \tag{A.11}$$

A marginal junior produces current flow surplus $\psi M - w_J$, separates at rate δ_J, and becomes a senior at rate λ. Promotion changes the firm's value from V_J to V_S, so the capital gain from promotion is $V_S - V_J$. The junior Bellman equation is therefore

$$rV_J = \psi M - w_J + \lambda(V_S - V_J) - \delta_J V_J. \tag{A.12}$$

Rearranging,

$$(r + \lambda + \delta_J)V_J = \psi M - w_J + \lambda V_S, \tag{A.13}$$

and hence

$$V_S = \frac{M - w_S}{B}, \qquad V_J = \frac{\psi M - w_J + \lambda V_S}{D}. \tag{A.14}$$

Using $M = AF_s(s, j)$ and $\psi M = A[\eta F_j(s, j) - \phi F_s(s, j)]$, these are the value functions stated in equation (3); the main-text notation corresponds to the normalisation $A = 1$.

The first-order condition for hiring with $C_i(h_i) = \frac{1}{2}\kappa_i h_i^2$ is

$$h_i = \frac{V_i}{\kappa_i}, \qquad i \in \{J, S\}. \tag{A.15}$$

Therefore

$$z^\star = \frac{h_J}{h_J + h_S} = \left[1 + \frac{\kappa_J}{\kappa_S}\frac{V_S}{V_J}\right]^{-1}. \tag{A.16}$$

Define the junior-to-senior value ratio

$$Q \equiv \frac{V_J}{V_S}. \tag{A.17}$$

For an interior steady state with $V_S > 0$ and $V_J > 0$, $Q > 0$, and

$$z^\star = \left[1 + \frac{\kappa_J}{\kappa_S} Q^{-1}\right]^{-1}. \tag{A.18}$$

Moreover,

$$\frac{\partial z^\star}{\partial Q} = \frac{(\kappa_J/\kappa_S)Q^{-2}}{\left[1 + (\kappa_J/\kappa_S)Q^{-1}\right]^2} > 0. \tag{A.19}$$

Signing $\partial z^\star/\partial\theta$ is therefore equivalent to signing $\partial Q/\partial\theta$.

Using $\mathrm{MP}_J = \psi M$, the value ratio is

$$Q = \frac{V_J}{V_S} = \frac{B(\psi M - w_J) + \lambda(M - w_S)}{D(M - w_S)}. \tag{A.20}$$

For any parameter $\theta \neq \lambda$ that affects Q only through M and ψ, with r, δ_i, w_i, and λ held fixed, differentiating (A.20) gives

$$\frac{\partial Q}{\partial\theta} = \frac{BM}{D(M - w_S)^2}\left[(M - w_S)\frac{\partial\psi}{\partial\theta} + (w_J - \psi w_S)\frac{1}{M}\frac{\partial M}{\partial\theta}\right]. \tag{A.21}$$

This master expression is useful for the supervision-cost and junior-productivity derivatives.

A.3 Learning-rate derivative

Because worker stocks are fixed, and because neither s nor j depends on λ, the production objects M and ψ are held fixed when differentiating with respect to λ. Differentiating (A.20) gives

$$\frac{\partial Q}{\partial \lambda} = \frac{1-Q}{r+\lambda+\delta_J}. \tag{A.22}$$

Since $z^\star$ is strictly increasing in Q,

$$\frac{\partial z^\star}{\partial \lambda} > 0 \quad \Longleftrightarrow \quad \frac{\partial Q}{\partial \lambda} > 0 \quad \Longleftrightarrow \quad Q < 1 \quad \Longleftrightarrow \quad V_S > V_J. \tag{A.23}$$

This proves (5).

It is also useful to record the current-surplus form of the capital-gain condition used in the discussion after the proposition. Let

$$\Pi_S \equiv M - w_S, \qquad \Pi_J \equiv \psi M - w_J. \tag{A.24}$$

Then $V_S > V_J$ is equivalent to

$$(r+\delta_J)\frac{\Pi_S}{r+\delta_S} > \Pi_J. \tag{A.25}$$

If $\Pi_J > 0$, this can be written as

$$\frac{\mathrm{MP}_S - w_S}{\mathrm{MP}_J - w_J} > \frac{r+\delta_S}{r+\delta_J}. \tag{A.26}$$

If $\Pi_J \le 0$, then (A.25) holds automatically whenever $V_S > 0$, so the ratio form (A.26) should be read as the positive-current-junior-surplus case.

A.4 Supervision-cost derivative

For the supervision-cost comparative static, worker stocks are fixed, so a change in ϕ changes effective senior time but not the numbers of senior or junior workers:

$$\frac{\partial s}{\partial \phi} = -\ell^J, \qquad \frac{\partial j}{\partial \phi} = 0. \tag{A.27}$$

Thus

$$\frac{\partial M}{\partial \phi} = -A\ell^J F_{ss}(s,j) > 0. \tag{A.28}$$

The inequality follows from homogeneity of degree one and strict concavity in the factor ratio. In particular, writing $F(s,j) = sf(j/s)$, strict concavity in the factor ratio implies $f'' < 0$, and hence $F_{ss} = (j/s)^2 f''(j/s)/s < 0$. The same representation gives $F_{sj} = -(j/s)f''(j/s)/s > 0$ and $F_{jj} = f''(j/s)/s < 0$.

Using $\psi = \eta F_j/F_s - \phi$,

$$\frac{\partial \psi}{\partial \phi} = -1 - \eta \ell^J \frac{F_{sj}F_s - F_j F_{ss}}{F_s^2} < 0. \tag{A.29}$$

The term in the numerator is positive: under the same homogeneity and strict-concavity conditions, $F_{sj} > 0$, $F_s > 0$, $F_j > 0$, and $F_{ss} < 0$, so $F_{sj}F_s - F_j F_{ss} > 0$.

Let $\mu_J \equiv w_J/\mathrm{MP}_J = w_J/(\psi M)$ and $\mu_S \equiv w_S/\mathrm{MP}_S = w_S/M$ denote the wage-to-marginal-product markdowns for juniors and seniors. The wage-share condition under which (6) holds is equivalent to the wage-wedge condition used in the proof: since $M > 0$ and $\psi > 0$,

$$\mu_J \le \mu_S \quad \Longleftrightarrow \quad \frac{w_J}{\psi M} \le \frac{w_S}{M} \quad \Longleftrightarrow \quad w_J \le \psi w_S. \tag{A.30}$$

Under (A.30), the second bracketed term in (A.21) is weakly negative because $\partial M/\partial \phi > 0$. The first bracketed term is strictly negative because $V_S > 0$ implies $M - w_S > 0$, while $\partial \psi/\partial \phi < 0$. The prefactor in (A.21) is strictly positive. Therefore

$$\frac{\partial Q}{\partial \phi} < 0. \tag{A.31}$$

Since $z^\star$ is strictly increasing in Q,

$$\frac{\partial z^\star}{\partial \phi} < 0. \tag{A.32}$$

This proves (6).

A.5 Junior productivity

The effect of junior productivity η is not globally signed in the general production-function environment. Again, the comparative static holds worker stocks fixed:

$$\frac{\partial s}{\partial \eta} = 0, \qquad \frac{\partial j}{\partial \eta} = \ell^J. \tag{A.33}$$

Therefore

$$\frac{\partial M}{\partial \eta} = A\ell^J F_{sj}(s,j) > 0. \tag{A.34}$$

The relative marginal-product term satisfies

$$\frac{\partial \psi}{\partial \eta} = \frac{F_j}{F_s} + \eta \ell^J \frac{F_{jj}F_s - F_j F_{sj}}{F_s^2}. \tag{A.35}$$

The first term is positive, while the second term is negative under homogeneity and strict concavity in the factor ratio. Hence $\partial\psi/\partial\eta$ is not globally signed.

This ambiguity can be stated more directly using the local elasticity of substitution. Let

$$\sigma(s,j) \equiv \left[\frac{\partial \ln(F_s/F_j)}{\partial \ln(j/s)}\right]^{-1} > 0. \tag{A.36}$$

Because s and ℓ^J are fixed when differentiating with respect to η, j/s moves one-for-one with η. Homogeneity implies

$$\frac{\partial \psi}{\partial \eta} = \frac{\eta F_j/F_s}{\eta}\left(1 - \frac{1}{\sigma(s,j)}\right) = \frac{\psi + \phi}{\eta}\left(1 - \frac{1}{\sigma(s,j)}\right). \tag{A.37}$$

Thus the direct relative-productivity term rises with η when $\sigma(s,j) > 1$, falls when $\sigma(s,j) < 1$, and is locally unchanged when $\sigma(s,j) = 1$. In addition, the master expression (A.21) contains the senior-value term $\partial M/\partial\eta > 0$, weighted by the wage wedge $w_J - \psi w_S$. Without additional restrictions on the local elasticity of substitution and the wage wedge, the sign of $\partial Q/\partial\eta$, and therefore the sign of $\partial z^\star/\partial\eta$, is ambiguous.

B General-Equilibrium Extension: Endogenous Wages

Appendix A treats wages as fixed when deriving the effect of WFH on the junior share of hiring. This appendix closes the model by allowing wages to clear type-specific hiring markets. The extension is deliberately minimal. The firm side is unchanged; the only new ingredient is an upward-sloping supply curve for each worker type. This lets wages adjust endogenously while preserving the main intuition of the fixed-wage model.

Write the supply of type $i \in \{J, S\}$ workers to the hiring market as

$$h_i^{\mathscr{S}} = H_i(w_i), \qquad H_i'(w_i) > 0.$$

A simple microfoundation is a representative household with quasi-linear consumption and separable disutility from supplying the two worker types to be hired by the representative firm:

$$U = c - \sum_{i \in \{J,S\}} \frac{\chi_i}{1 + 1/\varepsilon_i} \left(h_i^{\mathscr{S}}\right)^{1+1/\varepsilon_i}, \qquad \varepsilon_i > 0,\ \chi_i > 0, \tag{B.1}$$

with budget constraint

$$c = \Pi + \sum_{i \in \{J,S\}} w_i h_i^{\mathscr{S}},$$

where Π denotes firm profits rebated to the household.[26] Quasi-linearity and separability imply that profits and income do not feed back into labor-supply decisions. The household first-order conditions are

$$w_i = \chi_i \left(h_i^{\mathscr{S}}\right)^{1/\varepsilon_i}, \qquad \Rightarrow \qquad h_i^{\mathscr{S}} = H_i(w_i) = \left(\frac{w_i}{\chi_i}\right)^{\varepsilon_i}. \tag{B.2}$$

The derivations below use only $H_i'(w_i) > 0$, so the isoelastic form is not essential.

Normalize the mass of firms to one, so that firm hiring demand can be read as aggregate hiring demand. Let

$$R_S \equiv r + \delta_S, \qquad R_J \equiv r + \delta_J + \lambda.$$

It is useful to define the senior marginal product and the junior net static marginal product as

$$M(\phi) \equiv F_s(s, j), \qquad B(\phi) \equiv \eta F_j(s, j) - \phi F_s(s, j), \tag{B.3}$$

[26] In steady state $h_i^{\mathscr{S}}$ is the flow of new hires of type i that the household supplies per unit time, and w_i is reinterpreted as the per-hire annuity value of the indefinite employment contract; this keeps the budget constraint dimensionally consistent with the firm's continuous-time formulation. The comparative statics rely only on $H_i'(w_i) > 0$, so the choice of units in this microfoundation is innocuous.

where, as in the main text,

$$s = \ell_S - \phi \ell_J, \qquad j = \eta \ell_J.$$

Then the worker values are

$$V_S = \frac{M(\phi) - w_S}{R_S}, \qquad V_J = \frac{B(\phi) - w_J + \lambda V_S}{R_J}. \tag{B.4}$$

Firm demand for hires is still

$$h_i^{\mathscr{D}} = \frac{V_i}{\kappa_i}, \qquad i \in \{J, S\}. \tag{B.5}$$

An equilibrium is a pair of wages $(w_J^\star, w_S^\star)$ such that

$$H_S(w_S^\star) = \frac{V_S(w_S^\star; \phi)}{\kappa_S}, \tag{B.6}$$

$$H_J(w_J^\star) = \frac{V_J(w_J^\star, w_S^\star; \lambda, \phi)}{\kappa_J}. \tag{B.7}$$

The senior market solves first because V_S does not depend on w_J. Given $w_S^\star$, the junior market then clears. Under the maintained interiority assumptions $V_S > 0$ and $V_J > 0$, the equilibrium is unique: the supply side is increasing in the relevant wage, while the demand side is decreasing in the relevant wage.

The equilibrium junior hiring share is

$$z^{\star,GE} \equiv \frac{h_J}{h_J + h_S} = \left[1 + \frac{\kappa_J}{\kappa_S} \frac{V_S(w_S^\star)}{V_J(w_J^\star, w_S^\star)} \right]^{-1}. \tag{B.8}$$

This is the same expression as in the fixed-wage model, except that values are evaluated at equilibrium wages.

Useful incidence terms. Define the type-specific quantity pass-through terms

$$\tau_i \equiv \frac{H_i'(w_i)}{H_i'(w_i) + 1/(\kappa_i R_i)} \in (0, 1), \qquad i \in \{J, S\}, \tag{B.9}$$

where $R_i = R_J$ for juniors and $R_i = R_S$ for seniors. These terms measure how much of a fixed-wage demand shift shows up as an equilibrium quantity response. If labor supply is very elastic, $\tau_i \to 1$, and the model approaches the fixed-wage case. If labor supply is very inelastic, $\tau_i \to 0$, and the demand shift is absorbed mostly by wages.

Under the isoelastic supply curve in (B.2),

$$\tau_i = \frac{\varepsilon_i R_i V_i / w_i}{1 + \varepsilon_i R_i V_i / w_i} = \frac{\varepsilon_i R_i V_i}{w_i + \varepsilon_i R_i V_i}. \tag{B.10}$$

We now derive the comparative statics of $z^{\star,GE}$. To reduce notation, all objects are evaluated at the equilibrium and stars are suppressed.

For any parameter θ, start from

$$z^\star = \frac{h_J}{h_J + h_S}, \qquad 1 - z^\star = \frac{h_S}{h_J + h_S}.$$

Hence

$$\frac{z^\star}{1 - z^\star} = \frac{h_J}{h_S}.$$

Taking logs and differentiating gives

$$\frac{1}{z^\star(1 - z^\star)} \frac{dz^\star}{d\theta} = \frac{d \log h_J}{d\theta} - \frac{d \log h_S}{d\theta}.$$

Therefore,

$$\frac{dz^\star}{d\theta} = z^\star(1 - z^\star) \left[\frac{d \log h_J}{d\theta} - \frac{d \log h_S}{d\theta} \right]. \tag{B.11}$$

Thus the junior hiring share rises when junior hiring grows faster than senior hiring.

Effect of the learning rate λ. The senior value V_S does not depend directly on λ, so the senior wage and senior hiring are unchanged:

$$\frac{dw_S}{d\lambda} = 0, \qquad \frac{dh_S}{d\lambda} = 0. \tag{B.12}$$

For juniors,

$$V_J = \frac{B - w_J + \lambda V_S}{R_J}, \qquad R_J = r + \delta_J + \lambda.$$

The direct effect of λ on junior hiring demand, before the junior wage adjusts, is

$$\begin{aligned} \frac{\partial V_J}{\partial \lambda} &= \frac{V_S R_J - (B - w_J + \lambda V_S)}{R_J^2} \\ &= \frac{V_S R_J - R_J V_J}{R_J^2} = \frac{V_S - V_J}{R_J}. \end{aligned} \tag{B.13}$$

The total effect allows the junior wage to adjust. Differentiating the junior market-clearing condition

$$H_J(w_J) = \frac{V_J(w_J, w_S, \lambda)}{\kappa_J}$$

with respect to λ gives

$$H_J'(w_J)\frac{dw_J}{d\lambda} = \frac{1}{\kappa_J}\left[\frac{V_S - V_J}{R_J} - \frac{1}{R_J}\frac{dw_J}{d\lambda}\right].$$

Collecting terms,

$$\left[H_J'(w_J) + \frac{1}{\kappa_J R_J}\right]\frac{dw_J}{d\lambda} = \frac{1}{\kappa_J}\frac{V_S - V_J}{R_J}.$$

Hence

$$\frac{dw_J}{d\lambda} = \frac{(V_S - V_J)/R_J}{\kappa_J H_J'(w_J) + 1/R_J}. \tag{B.14}$$

Using $h_J = H_J(w_J) = V_J/\kappa_J$, the induced junior hiring response is

$$\begin{aligned}\frac{d\log h_J}{d\lambda} &= \frac{1}{V_J}\left[\frac{V_S - V_J}{R_J} - \frac{1}{R_J}\frac{dw_J}{d\lambda}\right] \\ &= \tau_J \frac{V_S - V_J}{R_J V_J}.\end{aligned} \tag{B.15}$$

Combining (B.11) and (B.15),

$$\frac{dz^{\star,GE}}{d\lambda} = z^\star(1 - z^\star)\tau_J \frac{V_S - V_J}{R_J V_J} \tag{B.16}$$

and therefore

$$\frac{dz^{\star,GE}}{d\lambda} > 0 \quad \Longleftrightarrow \quad V_S > V_J. \tag{B.17}$$

The condition is exactly the same as in the fixed-wage model. Endogenous wages only dampen the magnitude of the quantity response through τ_J. Intuitively, when a junior worker is an investment in future senior surplus, a higher promotion or learning rate raises the value of junior hiring. The junior wage rises endogenously, but with upward-sloping supply not all of the demand shift is absorbed by wages, so junior hiring also rises.

Effect of the supervision cost ϕ. The effect of ϕ works through two channels. First, a higher supervision cost directly lowers the net static contribution of a junior worker. Second, because each junior absorbs senior time, a higher ϕ reduces effective senior labor $s = \ell_S - \phi\ell_J$, which raises the marginal product of senior time.

The relevant fixed-wage derivatives are

$$V_{S\phi} = \frac{M_\phi}{R_S}, \qquad V_{J\phi} = \frac{B_\phi + \lambda V_{S\phi}}{R_J}. \tag{B.18}$$

Given $s = \ell_S - \phi\ell_J$, $j = \eta\ell_J$, and holding current employment stocks fixed when evaluating marginal products,

$$M_\phi = \frac{\partial F_s(s,j)}{\partial \phi} = -\ell_J F_{ss} > 0, \tag{B.19}$$

because $F_{ss} < 0$. Similarly,

$$\begin{aligned} B_\phi &= \frac{\partial}{\partial \phi}\left[\eta F_j(s,j) - \phi F_s(s,j)\right] \\ &= -\eta\ell_J F_{sj} - F_s + \phi\ell_J F_{ss} < 0. \end{aligned} \tag{B.20}$$

The inequality follows from $F_s > 0$, $F_{ss} < 0$, and $F_{sj} > 0$, where the last two properties follow from constant returns and strict concavity in the factor ratio.

The senior market-clearing condition is

$$H_S(w_S) = \frac{V_S(w_S;\phi)}{\kappa_S} = \frac{M(\phi) - w_S}{\kappa_S R_S}.$$

Differentiating with respect to ϕ gives

$$H_S'(w_S)\frac{dw_S}{d\phi} = \frac{1}{\kappa_S R_S}\left[M_\phi - \frac{dw_S}{d\phi}\right].$$

Collecting terms,

$$\left[H_S'(w_S) + \frac{1}{\kappa_S R_S}\right]\frac{dw_S}{d\phi} = \frac{M_\phi}{\kappa_S R_S}.$$

Thus

$$\frac{dw_S}{d\phi} = \frac{M_\phi}{1 + \kappa_S R_S H_S'(w_S)} > 0. \tag{B.21}$$

Thus a higher supervision cost raises the senior wage, because it raises the marginal product of effective senior time. The equilibrium change in the senior value is

$$\begin{aligned}\frac{dV_S}{d\phi} &= V_{S\phi} + V_{Sw_S}\frac{dw_S}{d\phi} \\ &= \frac{M_\phi}{R_S} - \frac{1}{R_S}\frac{M_\phi}{1+\kappa_S R_S H_S'(w_S)} \\ &= \tau_S V_{S\phi},\end{aligned} \tag{B.22}$$

and the senior hiring response is

$$\frac{d\log h_S}{d\phi} = \frac{1}{V_S}\frac{dV_S}{d\phi} = \tau_S \frac{V_{S\phi}}{V_S}. \tag{B.23}$$

For juniors, the direct effect of ϕ, allowing the senior market to clear but before the junior wage adjusts, is

$$\frac{B_\phi + \lambda\, dV_S/d\phi}{R_J} = \frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J}.$$

Differentiating the junior market-clearing condition with respect to ϕ gives

$$H_J'(w_J)\frac{dw_J}{d\phi} = \frac{1}{\kappa_J}\left[\frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J} - \frac{1}{R_J}\frac{dw_J}{d\phi}\right].$$

Collecting terms,

$$\left[H_J'(w_J) + \frac{1}{\kappa_J R_J}\right]\frac{dw_J}{d\phi} = \frac{1}{\kappa_J}\frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J}.$$

The induced junior hiring response is therefore

$$\begin{aligned}\frac{d\log h_J}{d\phi} &= \frac{1}{V_J}\left[\frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J} - \frac{1}{R_J}\frac{dw_J}{d\phi}\right] \\ &= \tau_J \frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J V_J}.\end{aligned} \tag{B.24}$$

Combining (B.11), (B.23), and (B.24),

$$\frac{dz^{\star,GE}}{d\phi} = z^\star(1-z^\star)\left[\tau_J\frac{B_\phi + \lambda\tau_S V_{S\phi}}{R_J V_J} - \tau_S\frac{V_{S\phi}}{V_S}\right]. \tag{B.25}$$

Therefore,

$$\frac{dz^{\star,GE}}{d\phi} < 0 \tag{B.26}$$

if and only if

$$\tau_J \frac{B_\phi + \lambda \tau_S V_{S\phi}}{R_J V_J} < \tau_S \frac{V_{S\phi}}{V_S}.$$

Since R_J, V_J, and V_S are positive, this is equivalent to

$$\tau_J \left(B_\phi + \lambda \tau_S V_{S\phi}\right) V_S < \tau_S V_{S\phi} R_J V_J. \tag{B.27}$$

Using $R_J V_J = B - w_J + \lambda V_S$, this condition becomes

$$\tau_J B_\phi V_S + \lambda \tau_J \tau_S V_{S\phi} V_S < \tau_S V_{S\phi}(B - w_J) + \lambda \tau_S V_{S\phi} V_S.$$

Rearranging gives

$$\tau_J B_\phi V_S < \tau_S V_{S\phi}(B - w_J) + \lambda \tau_S (1 - \tau_J) V_{S\phi} V_S. \tag{B.28}$$

This expression separates the two effects. The left-hand side is negative because $B_\phi < 0$: raising ϕ directly lowers the static value of junior labor. The first term on the right-hand side is non-negative whenever juniors are paid no more than their net static marginal product, $B \geq w_J$. The second term is also non-negative, since $0 < \tau_J < 1$. Hence a transparent sufficient condition for the WFH supervision-cost effect to reduce the junior hiring share is

$$B(\phi) \geq w_J. \tag{B.29}$$

That is, firms are not paying juniors more than their current net static marginal product before accounting for the promotion option.

The condition used in the fixed-wage proof in the main text is slightly stronger but immediately implies (B.29). Recall that

$$B(\phi) = \eta F_j - \phi F_s = \psi F_s, \qquad \psi \equiv \eta \frac{F_j}{F_s} - \phi.$$

If $w_J \leq \psi w_S$, then, since $V_S > 0$ implies $F_s > w_S$,

$$w_J \leq \psi w_S < \psi F_s = B(\phi).$$

Thus the same economically mild condition used in the fixed-wage model is sufficient for

$$\frac{dz^{\star,GE}}{d\phi} < 0. \tag{B.30}$$

Comparison with the fixed-wage model. The fixed-wage model is the limiting case in which both type-specific labor supplies are perfectly elastic, so $\tau_J \to 1$ and $\tau_S \to 1$. The endogenous-wage model preserves the main comparative statics.

First, the learning-rate result is unchanged in sign:

$$\frac{dz^{\star,GE}}{d\lambda} > 0 \quad \Longleftrightarrow \quad V_S > V_J.$$

Endogenous wages only scale the fixed-wage effect by τ_J. A higher learning rate raises junior demand when juniors are valuable as future seniors; the junior wage absorbs part of this demand shift, but junior hiring still rises when supply is upward sloping.

Second, the supervision-cost result is also unchanged under the same sufficient condition used in the fixed-wage model. A higher ϕ lowers the junior net static marginal product and raises the marginal product of senior time. With endogenous wages, part of the senior demand shift is absorbed by a higher senior wage, and part of the junior demand shift is absorbed by a change in the junior wage. These wage adjustments affect magnitudes, but they do not overturn the result that a higher supervision cost lowers the junior share of hiring when juniors are not being paid above their net static marginal product.

The extension also gives additional wage predictions. A rise in ϕ raises the senior wage:

$$\frac{dw_S}{d\phi} > 0.$$

The junior wage response to ϕ is

$$\frac{dw_J}{d\phi} = \frac{(B_\phi + \lambda\tau_S V_{S\phi})/R_J}{\kappa_J H_J'(w_J) + 1/R_J}, \tag{B.31}$$

so the junior wage falls when the direct loss in junior static productivity dominates the attenuated continuation-value gain:

$$B_\phi + \lambda\tau_S V_{S\phi} < 0.$$

For the learning-rate margin, the senior wage is unchanged:

$$\frac{dw_S}{d\lambda} = 0.$$

The junior wage response is

$$\frac{dw_J}{d\lambda} = \frac{(V_S - V_J)/R_J}{\kappa_J H_J'(w_J) + 1/R_J}.$$

Thus the junior wage rises with λ if and only if $V_S > V_J$. Equivalently, a WFH-induced fall in the learning rate lowers the junior wage whenever junior hiring is an investment in future senior surplus. In this sense, the endogenous-wage model yields the same qualitative WFH prediction as the fixed-wage model: WFH lowers the junior share of hiring by raising supervision costs and slowing on-the-job learning, while equilibrium wages absorb part of the shock.

C Additional Remarks on Data

C.1 Sample Overview by Country

Tables C.1 and C.2 provide summary statistics for the basic size and structure of the two microdata sources used throughout this paper. These are the Revelio Labs data on new hires, which contain 243 million new employer–employee matches from 9.9 million unique firms across 722 6-digit O*NET occupations over 2017–2025. Our data on job vacancy postings comes from Lightcast, which contains 407 million job postings from 3.7 million unique employers across 702 occupations over the same window. Table C.3 further reports the country-level breakdown of new-hires by seniority level. In our main paper, we combine the bottom two levels into one 'junior' category. Here, we separately report both the bottom category (entry-level), where monthly hiring has fallen by 14–29% from 2017-19 to 2023-25, the junior category, which saw declines of around 12-26%, and senior, which has risen by between 5-21%.

C.2 Missingness in Experience Requirements (Lightcast)

A non-trivial share of Lightcast postings do not include a pre-extracted value for the years experience covariate. This is not unusual, because a great many postings simply omit this information when they do not have such requirements.[27] Nonetheless, we do not set all missing values to zero, we instead use an inverse-probability-weighting correction. For each 3-digit-occupation×year stratum, we fit a fixed-effect Poisson model with the experience-reporting binary indicator as the outcome, log posting length (in characters) as the regressor, and fixed effects for detailed (6-digit) occupation, posting source (i.e. which website the posting was scraped from), and year×month. This yields, for each posting i, a predicted probability of report-

[27] The share of postings without an experience requirement is roughly 54–55% in the United States, 75% in the United Kingdom, 60% in Canada, and 80% in Australia, varying modestly across the 2017–2019 and 2023–2025 windows. This level of missingness is expected, as it is very often (but not always) the case that jobs requiring no experience simply omit this information from the posting text.

ing $\hat{\pi}_i = \Pr(\text{experience reported} \mid X_i)$. We then reweight observed postings by $1/\hat{\pi}_i$, normalised within stratum so that the per-stratum sum of weights equals the observed count. These weights are then used whenever we aggregate postings to the cell level, and ensures we are using a selection-adjusted distribution of experience requirements.

C.3 Firm Matching

In our complimentary analysis of actual WFH adoption, we measure this using Lightcast job postings. We then need to match this to Revelio Labs hiring, which is done by matching employer names across these two sources. analysis requires linking employers across two sources that share no common identifier: Lightcast postings use self-reported employer names, while Revelio Labs assigns each firm an internal identifier. We construct the link from firm names alone, before using any hiring outcomes.[28] Each name is embedded as a high-dimensional text vector using OpenAI's `text-embedding-3-large` model — the strongest of OpenAI's general-purpose embedding endpoints — and candidate links are generated by bidirectional nearest-neighbour search within country-by-state blocks.[29] We tune the cosine threshold using a language-model-adjudicated learning sample of roughly 10,000 candidate pairs that classifies each pair as same firm, different firms, or ambiguous.[30] At $\theta^\star = 0.85$ the matched-pair classifier attains precision 0.78, recall 0.86, and $F_1 = 0.82$ on the adjudicated sample, with the positive class defined as same-firm: 78% of pairs above $\theta^\star$ are true same-firm matches, and the matching procedure captures 86% of the true same-firm pairs in the adjudicated sample. Applying the selected threshold yields 7,928,899 matched firm pairs, linking 3,771,033 unique Lightcast employer strings to 3,471,705 unique Revelio Labs employers. Because the link is many-to-many, we aggregate Lightcast firm-level actual-WFH adoption to the Revelio Labs employer using posting-count weights before estimating firm-level specifications.

[28] Raw names are retained as posted, augmented with Revelio subsidiary and parent-name fields, and deduplicated to 17,142,374 unique strings.

[29] We use a top-5 nearest-neighbour search with a cosine floor of 0.70, which yields 73,847,609 candidate Lightcast–Revelio Labs firm pairs across 91 country-by-state blocks.

[30] Ambiguous pairs are excluded from threshold choice; the selected threshold $\theta^\star = 0.85$ maximises F_1 on the adjudicated sample.

Table C.1: New-hires Data: Sample Overview

Country	New Hires	Firms	States/Regions	Occupations
United States	166,491,524	5,989,773	51	722
United Kingdom	38,743,924	1,990,651	4	722
Australia	15,765,880	853,187	8	722
Canada	22,429,542	1,112,501	13	722
Pooled	243,430,870	9,946,112	76	722

Notes: Data source: new employer–employee matches assembled from professional resumes, 2017–2025 (see Section 4). New Hires = total newly formed employer–employee matches across the full sample period. Firms = unique employer identifiers. States/Regions = unique geographic units (US: states; UK: nations; Canada: provinces; Australia: states/territories). Occupations = unique 6-digit O*NET 2019 codes. Five 2-digit O*NET groups (21, 25, 33, 45, 55) are excluded from the analysis sample.

Table C.2: Job-postings Data: Sample Overview

Country	Job Postings	Firms	Regions	Occupations
United States	318,529,861	2,569,931	51	702
United Kingdom	56,791,794	595,363	391	631
Australia	9,901,723	181,772	88	631
Canada	21,862,295	401,592	13	631
Pooled	407,085,674	3,748,658	543	702

Notes: Data source: online job vacancy postings collected from job boards, recruitment websites, and employer websites, 2017–2025 (see Section 4). Postings = total job vacancy postings across the full sample period. Firms = unique employer names. Regions = unique geographic units (US: states; UK: LAU1 areas; Canada: provinces; Australia: SA4 areas). Occupations = unique 6-digit O*NET 2019 codes (non-US mapped via LOT→O*NET crosswalk). Five 2-digit O*NET groups (21, 25, 33, 45, 55) are excluded from the analysis sample.

Table C.3: Hiring by Seniority-level and Country in the New-hires Data

		Avg. Monthly Matches		Growth (%)		
Country	Seniority	2017-19	2023-25	Raw	Within-Occ	Within-State
United States	Entry Level	731,769	519,385	-29.0	-25.3	-28.9
	Junior	404,812	301,329	-25.6	-24.0	-25.5
	Senior	544,640	573,038	5.2	5.8	5.0
United Kingdom	Entry Level	140,269	108,663	-22.5	-18.1	-21.9
	Junior	99,244	77,580	-21.8	-21.3	-21.9
	Senior	137,152	158,078	15.3	18.5	14.6
Australia	Entry Level	57,830	44,012	-23.9	-20.6	-24.3
	Junior	40,737	32,074	-21.3	-22.0	-21.2
	Senior	55,500	62,493	12.6	15.2	13.2
Canada	Entry Level	96,387	82,581	-14.3	-11.3	-14.3
	Junior	52,120	45,727	-12.3	-10.2	-12.3
	Senior	62,024	74,781	20.6	22.4	20.2
Pooled	Entry Level	1,026,256	754,641	-26.5	-22.8	-26.2
	Junior	596,913	456,710	-23.5	-22.3	-23.5
	Senior	799,315	868,391	8.6	9.9	8.3

Notes: Data source: new employer–employee matches assembled from professional resumes (see Section 4). Seniority levels assigned by the data provider based on job title parsing. Raw growth compares average monthly matches across periods. Within-Occupation growth holds the 2019 US 3-digit ONET distribution fixed. Within-State growth holds each country's 2019 state/region distribution fixed.

D Additional Remarks on Motivating Evidence

D.1 AI and WFH Adoption: External Corroborating Series

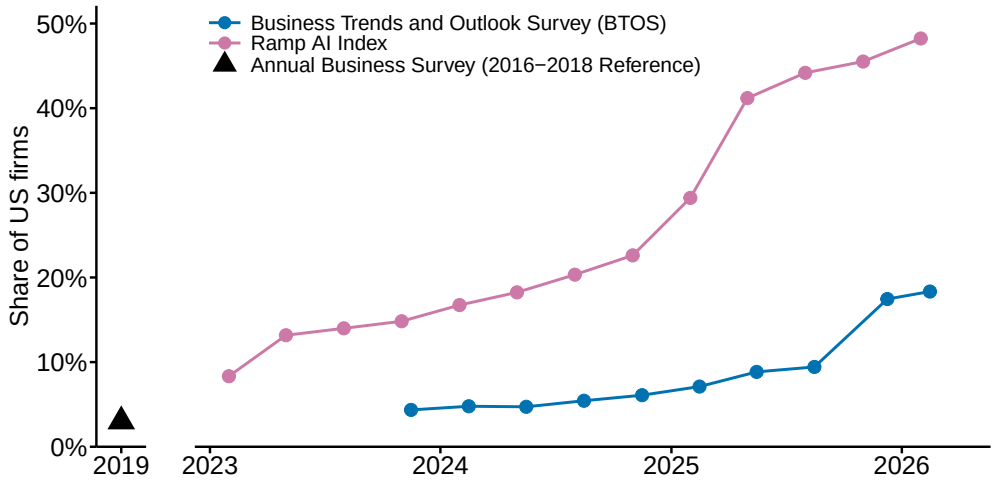
The post-2022 break in junior hiring coincides with sharply rising generative-AI adoption and with the persistence of pandemic-era remote-work levels at several multiples of the pre-2020 baseline. Figure D.1 plots contemporaneous aggregate adoption paths for these two technologies. Panel (a) combines two US firm-level AI-use series: the Census Bureau's Business Trends and Outlook Survey (BTOS) and the Ramp AI Index (Kharazian, 2026). Both rise from a near-zero base in early 2023, with BTOS roughly tripling and the Ramp index rising about six-fold over the subsequent two years. The 2019 McElheran et al. (2024) estimate from the Annual Business Survey, which predates ChatGPT, sits at about 3 percent and anchors the pre-period. Panel (b) plots the share of online job postings mentioning remote or hybrid work in our main-analysis countries, using the index of Hansen et al. (2023) distributed at WFHmap.com. Remote-work mention shares rose roughly four-fold during the pandemic, remain at multiples of their 2019 baseline in each country, and are broadly stable in the post-2022 window during which we estimate our DiD specifications.

D.2 Cross-Country Summary and Descriptive Facts

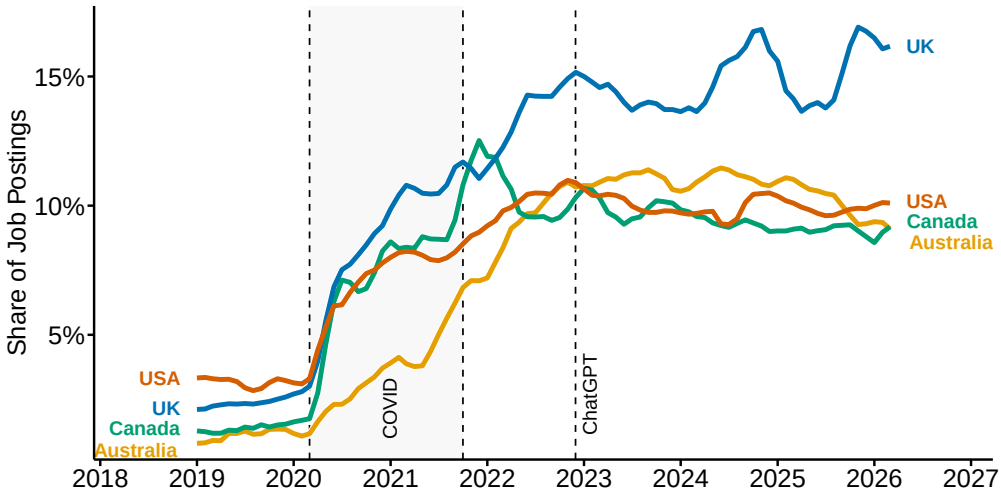
Figure D.2 decomposes the US decline into absolute flows by seniority level.

Figure D.1: AI Adoption Accelerated in 2025 While WFH Stabilised Above Pre-Pandemic Levels

(a) US firm-level AI adoption: BTOS, Ramp, and 2019 ABS benchmark

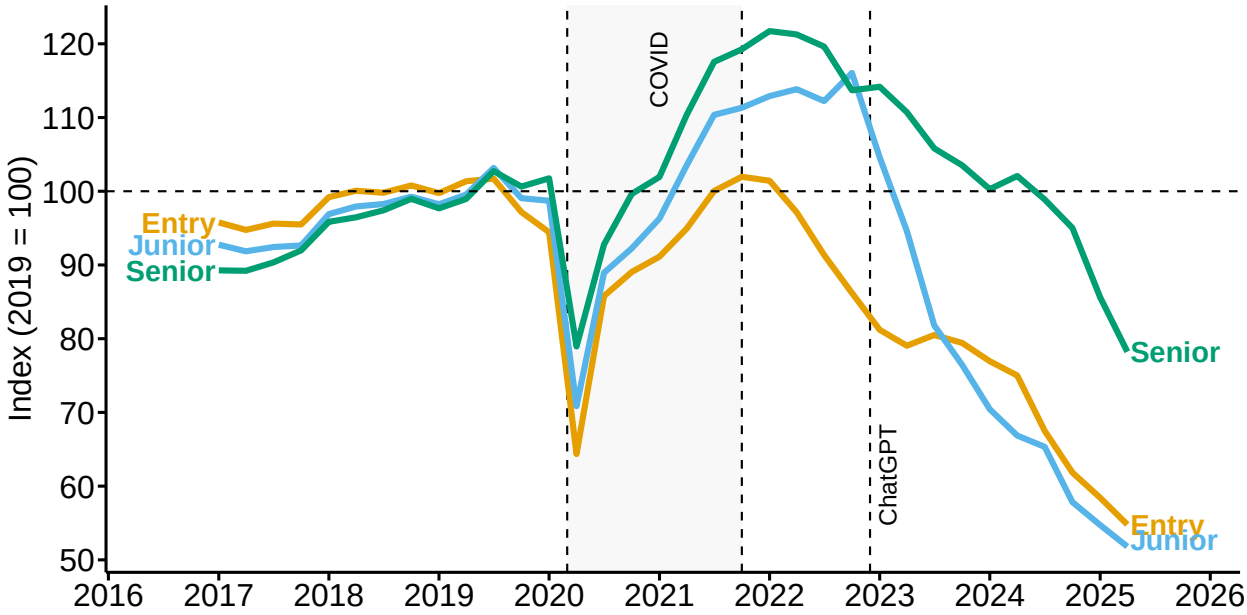


(b) Share of postings mentioning remote/hybrid work (3-month MA)



Note: *Panel (a)* plots three US firm-level AI adoption series. *BTOS (navy):* the Census Bureau's biweekly Business Trends and Outlook Survey, covering roughly one million US firms. Until 2025-11-16 it asked about AI use "in producing goods or services"; the two right-most quarterly points use a broader wording ("in any of its business functions") that captures uses outside production. *Ramp (pink):* Kharazian (2026) reports the monthly share of ~40,000 US Ramp corporate-card customers with any spend on an AI vendor (OpenAI, Anthropic, Microsoft Copilot, Perplexity, etc.), collapsed to quarterly means. *2019 ABS (black triangle):* McElheran et al. (2024)'s tabulation of the Annual Business Survey question on AI use in 2016–2018; the x-axis is broken to separate this pre-ChatGPT benchmark from the 2023–2026 series. *Panel (b)* plots the monthly share of online job vacancies mentioning remote or hybrid work in each country, as a 3-month moving average, from Hansen et al. (2023) (WFHmap.com). Dashed vertical lines mark the onset of COVID (2020-03-01), the end of broad COVID restrictions (2021-10-01), and the release of ChatGPT (2022-11-30); the shaded grey rectangle in Panel (b) covers the COVID window.

Figure D.2: Index of New Hires by Seniority Level (USA, 2019 = 100)



Note: Quarterly series of US new employer–employee matches from Revelio Labs, 2017–2025, broken down by seniority. "Entry" and "Junior" use the bottom-two groups of Revelio's six-level seniority scale, which combines job description with each worker's past experience at the start of their tenure (see Section 4); "Senior" aggregates the remaining four levels. Within each 3-digit O*NET occupation × seniority cell, raw counts are scaled by the cell's 2019 average and then aggregated using each occupation's 2019 share of total US new hires (common weights across seniority levels), so the index removes occupation-mix composition effects. Each series is seasonally adjusted using the U.S. Census Bureau's X-13ARIMA-SEATS procedure and re-indexed so its 2019 mean equals 100.

E Additional Empirical Results

E.1 Using the Senior Wage Premium and Replacement Rate as additional outcomes.

We report results using two additional outcomes of interest, the senior-to-junior wage premium and the (relative) replacement rate of junior workers to senior workers. This helps to re-enforce the implications of our conceptual framework and general equilibrium extension, which predicts a rise in supervision cost ϕ, or a fall in the learning rate λ, should both raise the senior–junior starting-wage gap and lower the rate at which junior hires keep pace with junior separations (relative to seniors). Table E.1 reports these results, but we first define explicitly the measurement of both outcomes.

The senior wage premium and the replacement rate are

$$\omega_{it} \;=\; 100 \times \left(\overline{\log w}^{S}_{it} - \overline{\log w}^{J}_{it}\right), \tag{E.1}$$

$$r^{k}_{it} \;=\; \frac{H^{k}_{it}}{D^{k}_{it}}, \qquad k \in \{J, S\}, \tag{E.2}$$

where H^{k}_{it} and D^{k}_{it} denote new hires and separations of seniority k in unit i at time t. The *relative* junior replacement rate is then defined as the log ratio of the junior to the senior replacement rate,

$$\rho_{it} \;=\; 100 \times \log\!\left(\frac{r^{J}_{it}}{r^{S}_{it}}\right) \;=\; 100 \times \left[\log\!\left(\frac{H^{J}_{it}}{D^{J}_{it}}\right) - \log\!\left(\frac{H^{S}_{it}}{D^{S}_{it}}\right)\right], \tag{E.3}$$

so $\rho_{ut} = 0$ corresponds to junior and senior workforces replenishing at equal rates, and $\rho_{ut} < 0$ to net erosion of the junior workforce relative to senior. Salaries entering (E.1) are partially imputed by our data provider and should be read as a proxy signal rather than as real posted wages.

Table E.1 shows that both of these additional margins move in the direction predicted by our simple framework. In the joint specification of Table E.1, a one-standard-deviation increase in WFH exposure raises ω_{it} by 6.08 log points in Panel A ($p < 0.01$) and 7.20 log points in Panel B ($p < 0.01$), and lowers ρ_{it} by 2.29 ($p < 0.05$) and 4.21 log points ($p < 0.01$) respectively. The GenAI coefficients in the same regressions are small, insignificant, or signed incorrectly relative to the hypothesis that GenAI reduced demand for junior workers. This evidence further points to a period where wages and hiring are adjusting to reflect weaker demand for junior workers relative

to senior workers, and that this adjustment is more strongly associated with WFH exposure than with GenAI exposure.

E.2 Cross-Country Results

We next assess our main results, run separately across each of our four countries. Table E.2 reports these estimates from the joint specification of (9) for each country, with Panel A reporting results for the junior-share of new hires, and Panel B for the share of new online job vacancy postings requiring no more than 3 years experience from applicants.

In Table E.2, the WFH exposure coefficient in Panel A sits in a fairly narrow band of −1.28 to −1.70 at the state×occupation level (columns 1–4) and −2.44 to −2.74 percentage points at the firm×2-digit-occupation level (columns 5–8), with $p < 0.01$ in all eight specifications. This cross-country agreement underwrites our decision to pool the data in our main approach. The GenAI coefficient in the same eight regressions is insignificant in seven cases-the only exception is the UK state×occupation estimate of -0.92^{***}.

In Table E.2 Panel B, the results are less uniform across countries. The WFH coefficient is large and significant in the US (-1.48^{***} occ, -1.72^{***} firm) and Australia (-1.11^{***} occ, -0.95^{***} firm), and statistically null at the region-by-occupation level in the UK (−0.22) and Canada (−0.18). At the firm-by-narrow occupation level, the effects in Canada are significant (-1.01^{***}) while the UK null persists (−0.18). The country-specific GenAI coefficients on the posting outcome are positive-signed in the US and Canada, echoing the result in our main pooled designs. Overall, the cross-country results in Table E.2 are broadly consistent with the main pooled results, and suggest that the WFH exposure effect is more stable across countries than the GenAI exposure effect.

E.3 Robustness to 2-Digit Occupation Composition

A natural concern is that the joint WFH + GenAI specification is identified off a single 2-digit occupation cluster — for example, computer and mathematical occupations (SOC 15) or office and administrative support (SOC 43), where both exposures and the analysis sample are heavily concentrated. Table E.3 re-estimates the joint additive specification of columns (3) and (7) in Table 1, dropping each of the 18 included SOC 2-digit major groups in turn. Each row reports the eight WFH and GenAI coefficients of interest under the four designs (Panel A new hires, Panel A postings, Panel B new hires, Panel B postings) after dropping all data for the focal 2-digit group.

The pattern documented in the main text — a stable, negative, and statistically significant WFH coefficient alongside a sharply attenuated GenAI coefficient — holds across the leave-one-out replications.

E.4 Robustness to Alternative WFH and GenAI Measures

Table E.4 re-estimates the specification of columns (4) and (8) of Table 1, using alternative measures of WFH and GenAI exposure. This includes the teleworkability index of Dingel and Neiman (2020), the exposure to language models from Felten et al. (2023), and the Anthropic Economic Index of Appel et al. (2025). Throughout, the coefficient on WFH exposure is negative and statistically significant in all 24 cells of the table, and is larger in magnitude than the GenAI exposure coefficient when using our preferred WFH exposure measure. The effects of GenAI are significant and negatively signed in seven of twenty-four combinations, but never counteract the negative and significant WFH coefficient in the same cell. The main outlier is the penultimate row of Column (1) and (2), which show a stronger effect for GenAI exposure than WFH exposure. This may be driven by the fact that the teleworkable index of Dingel and Neiman (2020) is a binary measure at the occupation-level, rather than an intensity-based measure. It also may relate to the fact that the Felten et al. (2023) has the highest correlation to our preferred WFH measure, running into issues of multi-collinearity in the joint specification. The interaction terms appear mostly insignificant, with a few exceptions, suggesting that the effects of WFH and GenAI exposure are mostly additive rather than multiplicative in our data. Overall, this table supports our main claim that WFH exposure is a robust and consistent predictor of the decline in junior hiring, and can act as an important omitted variable when testing the effects of GenAI exposure on early career hiring outcomes.

E.5 Residualized Exposure

Here we discuss an alternative design which uses single-treatments, after residualizing each on the co-treatment. Specifically, we take a focal exposure Z^k and regress it on the other exposure Z^ℓ with the same fixed effects as in the main specification, and then re-standardise the residuals:

$$\widetilde{Z}_j^k = \text{std}\Big(Z_j^k - \widehat{a}^k - \widehat{\pi}^k Z_j^\ell\Big), \tag{E.4}$$

We then replace Z_j^k with $\widetilde{Z}_j^k$ in the region-by-narrow-occupation single-treatment design, or else use this to construct the shift-share exposure for the firm-by-broad occupation design.

Table E.5 reports the results from this single-treatment difference-in-differences using the residualized exposures. Columns (1) and (2) report the junior-share of new hires outcome, cols (3) and (4) the share of job postings requiring no more than 3 years experience. The same fixed effects and inference decisions from our main specification are used. These results also recover the asymmetric pattern, where (residualized) WFH-exposure is significant and negative, but residualized GenAI exposure is not.

E.6 Alternative Thresholds and Outcomes

Alternative functional forms and specifications. Our headline empirical design uses a linear OLS functional form with the shares (in percentage points) as the outcomes of interest. This is attractive for three reasons. First, the linearity of the share yields directly interpretable coefficients, i.e. the percentage-point changes per standard deviation of exposure. Second, with two-way fixed effects and millions of weighted cells, OLS is computationally light and does not introduce convergence or incidental-parameter concerns relative to, say, using granular fixed effects in a fractional-logit specification. Third, the residuals from the joint OLS specification are unimodal and centred at zero in both data sources (even though our cluster-robust standard errors are valid asymptotically without the Gaussian assumption). We nonetheless show alternative specification choices common in these settings in Table E.7.

The first alternative specification in Table E.7 panel A is an unweighted OLS design, removing all weights. Panel B reports the marginal effects associated with the fractional Poisson model of Papke and Wooldridge (1996). Panel C replaces the share measure with the log-odds transform. To keep the three estimators directly comparable, every coefficient in Table E.7 is reported as the implied marginal effect on the junior share, in percentage points, evaluated at the cell-weighted mean share $\bar{s}$, with delta-method standard errors.

Across these alternative designs, eleven out of twelve WFH coefficients across the three alternative specifications are negatively signed and statistically significant at conventional levels. The lone exception is the LinkedIn occupation-level estimate under the log-odds specification (Panel C), which is negatively signed but not significant. In the majority of cells, the implied response gradient is at least as strong as our headline specification, and the fractional Poisson (Panel B) delivers uniformly stronger implied responses.

This is not surprising, as all these specifications appear to deliver a similarly good approximation of the conditional mean of interest. Figure E.1 plots the kernel densities of the raw share outcomes and shows that even before our highly saturating fixed effects are applied, there is

sizable variation in these shares. Once the granular fixed effects are applied to recentre each units share-values, the identifying variation looks roughly normally distributed across both outcomes and both levels of aggregation, as shown in Figure E.2. This normality is not required for consistent and unbiased inference, but it does further justify the use of OLS on the shares.

Finally, we consider alternative threshold to define the share of new job postings going to early-career workers. Table E.6 reports the results pertaining to Equation (9) under alternative experience-threshold definitions of "junior". This includes the share of postings that require at most one year of experience, through to at most six. This highlights an invariance to the specific threshold number of years we choose to define 'junior' recruitment.

Table E.1: WFH/GenAI Exposure Effects on the Senior-Wage Premia and the Relative Replacement Rate

	Log Wage Premium: Starting Salary log(Senior) − log(Junior) (×100)				Log Relative Junior Replacement Rate (×100)			
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: State × Occ. × Month								
Post × WFH	5.76***		6.08***	6.82***	-1.808***		-2.293**	-2.653**
	(0.83)		(1.18)	(1.11)	(0.555)		(0.963)	(1.020)
Post × GenAI		3.41***	-0.51	-2.34**		-0.739**	0.754	1.637
		(0.79)	(1.08)	(1.09)		(0.337)	(0.776)	(1.070)
Post × WFH × GenAI				-2.88***				1.369**
				(0.60)				(0.659)
State × Occ. FE	✓	✓	✓	✓	✓	✓	✓	✓
State × Month FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	1,259,935 Cells / 115,825,863 Hires				811,166 Cells / 113,558,714 Hires			
R^2	0.56	0.55	0.56	0.56	0.40	0.39	0.40	0.40
Panel B: Firm × 2-Digit Occ. × Year								
Post × WFH	5.86***		7.20***	7.91***	-3.231***		-4.213***	-4.210***
	(0.12)		(0.21)	(0.23)	(0.091)		(0.138)	(0.165)
Post × GenAI		4.81***	-1.55***	-2.17***		-2.569***	1.149***	1.145***
		(0.13)	(0.22)	(0.23)		(0.079)	(0.103)	(0.155)
Post × WFH × GenAI				-0.59***				-0.003
				(0.10)				(0.078)
Firm × Occ. FE	✓	✓	✓	✓	✓	✓	✓	✓
Country × Year FE	✓	✓	✓	✓	✓	✓	✓	✓
Observations	3,300,042 Cells / 88,413,784 Hires				511,273 Cells / 52,850,960 Hires			
R^2	0.52	0.52	0.52	0.52	0.25	0.25	0.25	0.25

Notes: New-hires data. *Columns (1)–(4)*: Outcome is the senior wage premium ω_{ut} defined in equation (E.1), the difference between mean log starting salary of senior hires and that of entry/junior hires within each cell, ×100. Both mean log salary levels and the log premium are winsorized at the 1st and 99th percentiles. *Columns (5)–(8)*: Outcome is the log relative junior replacement rate ρ_{ut} defined in equation (E.2): 100 × [log(Junior share of hires) − log(Junior share of separations)]. A value of zero implies the junior share of the workforce is stable; negative values indicate net erosion. Cells require ≥ 10 total hires, ≥ 10 total separations, and ≥ 5 junior separations. *Panel A*: Unit is State×Occupation×Month. "State" denotes Province in Canada and Devolved Nation in the UK. Standard errors two-way clustered by ONET and State. Weighted by total hires. *Panel B*: Unit is Firm×2-Digit-Occupation×Year. Each firm×occupation cell must appear in at least one pre-period year (2017–19) and one post-period year (2023–25). Exposure is the 2-digit ONET major-group average of 7-digit occupation-level exposure. Standard errors clustered by firm. Outcome and exposures winsorized at the 1st and 99th percentiles. Exposures standardized. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

Table E.2: Country-by-Country Results: Joint WFH + GenAI Specification

	State × 6-Digit Occupation × Month Cells				Firm × 2-Digit Occupation × Year Cells			
	US	UK	Australia	Canada	US	UK	Australia	Canada
	(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)
Panel A: Dep. Var. Junior Share of New Hires								
Post × WFH (1sd)	-1.40***	-1.28***	-1.66***	-1.70***	-2.44***	-2.58***	-2.74***	-2.65***
	(0.39)	(0.20)	(0.41)	(0.31)	(0.75)	(0.70)	(0.80)	(0.69)
Post × GenAI (1sd)	-0.35	-0.92***	-0.40	-0.08	-0.11	-0.16	0.33	0.34
	(0.37)	(0.23)	(0.35)	(0.28)	(0.72)	(0.72)	(0.73)	(0.64)
Partial F-Statistic	13.0/0.9	41.9/16.0	16.5/1.3	29.7/0.1	10.7/0.0	13.6/0.1	11.7/0.2	14.9/0.3
Variance Inflation Factor	1.8	1.7	1.7	1.8	2.8	2.6	2.6	2.7
Condition number	2.2	2.2	2.1	2.2	3.0	2.9	2.9	3.0
Observations	1,663,257	134,136	225,016	289,777	3,677,092	798,663	349,245	434,144
Panel B: Dep. Var. Share of Postings Requiring ≤ 3 Years Experience								
Post × WFH (1sd)	-1.48***	-0.22	-1.11***	-0.18	-1.72***	-0.18	-0.95***	-1.01***
	(0.22)	(0.19)	(0.25)	(0.18)	(0.21)	(0.20)	(0.16)	(0.14)
Post × GenAI (1sd)	0.33**	-0.10	0.18	0.28**	0.47**	-0.06	0.11	0.30**
	(0.14)	(0.18)	(0.16)	(0.14)	(0.21)	(0.15)	(0.09)	(0.13)
Partial F-Statistic	45.3/5.4	1.3/0.3	19.6/1.4	0.9/4.1	65.7/5.1	0.8/0.2	33.1/1.5	51.0/4.9
Variance Inflation Factor	1.8	1.8	1.7	1.8	3.1	3.4	3.1	2.8
Condition number	2.2	2.2	2.1	2.3	3.2	3.4	3.2	3.0
Observations	1,829,657	3,837,332	635,973	274,845	1,355,270	995,413	180,893	457,925

Notes: This table reports the joint additive WFH + GenAI specification of Table 1, estimated separately for each country. The sample, outcome data, exposures, fixed effects, weights, and clustering are identical to Table 1 (within-country, since each regression includes only one country). Cols. (1)–(4) match the State×6-digit-Occupation×Month occ-level specification of Table 1 Panel A col. 3, with State×Occupation and State×Month fixed effects, weighted by total hires/postings, two-way clustered by Occupation and State. Cols. (5)–(8) match the Firm×2-Digit-Occupation×Year firm-level specification of Table 1 Panel B col. 3, here estimated at the unit level (rather than the BHJ shock level used in the pooled Panel B), with Firm×Occupation and Year fixed effects, weighted by total hires/postings, two-way clustered by firm and 2-digit occupation. Exposures are standardised on the pooled four-country sample. The Observations row reports the unit×period count entering each regression (after singleton-FE/NA filtering). The variance inflation factor (VIF) and condition number reported in each column are computed from the regression hessian after fixed-effect projection and weighting, restricted to the focal regressors actually present in the spec, so they reflect the residual variation that identifies each coefficient. Partial F is the squared t-ratio of each focal coefficient (reported as WFH/GenAI). Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.3: Main Results are Robust to Dropping any single 2-digit ONET-SOC Occupation Group

		Panel A: State×Occ×Month				Panel B: Firm×2d-Occ×Year			
		New Hires		Postings ≤ 3 yrs		New Hires		Postings ≤ 3 yrs	
Code	Occupation Group Name	WFH (1)	GenAI (2)	WFH (3)	GenAI (4)	WFH (5)	GenAI (6)	WFH (7)	GenAI (8)
11	Management	-1.59***	-0.37	-1.12***	0.25**	-1.90***	-0.31	-1.20***	0.11
13	Business and Financial Operations	-1.46***	-0.27	-1.48***	0.33***	-1.64***	-0.39	-1.22***	0.10
15	Computer and Mathematical	-1.40***	-0.39	-1.20***	0.26**	-1.39***	-0.18	-1.07***	0.10
17	Architecture and Engineering	-1.43***	-0.42	-1.24***	0.32***	-1.57***	-0.46	-1.15***	0.07
19	Life, Physical, and Social Science	-1.34***	-0.57*	-1.24***	0.27**	-1.52***	-0.54	-1.16***	0.07
23	Legal	-1.23***	-0.57*	-1.25***	0.27**	-1.50***	-0.50	-1.19***	0.09
27	Arts, Design, Entertainment, Sports, and Media	-1.54***	-0.48	-1.22***	0.24**	-1.64***	-0.51	-1.16***	0.06
29	Healthcare Practitioners and Technical	-1.54***	-0.37	-1.19***	0.24**	-1.70***	-0.42	-0.95***	-0.06
31	Healthcare Support	-1.38***	-0.43	-1.24***	0.27**	-1.53***	-0.43	-1.15***	0.07
35	Food Preparation and Serving Related	-1.41***	-0.29	-1.24***	0.28**	-1.54***	-0.35	-1.15***	0.07
37	Building and Grounds Cleaning and Maintenance	-1.43***	-0.40	-1.25***	0.28**	-1.57***	-0.44	-1.14***	0.05
39	Personal Care and Service	-1.44***	-0.41	-1.24***	0.25**	-1.56***	-0.46	-1.14***	0.05
41	Sales and Related	-1.50***	-0.17	-1.18***	0.22*	-1.62***	-0.39	-1.15***	0.05
43	Office and Administrative Support	-1.11**	-0.88*	-1.24***	0.19	-1.14***	-1.03***	-1.20***	0.07
47	Construction and Extraction	-1.45***	-0.35	-1.26***	0.34***	-1.57***	-0.44	-1.14***	0.05
49	Installation, Maintenance, and Repair	-1.42***	-0.35	-1.25***	0.27**	-1.56***	-0.42	-1.13***	0.02
51	Production	-1.44***	-0.35	-1.25***	0.28**	-1.56***	-0.42	-1.14***	0.06
53	Transportation and Material Moving	-1.42***	-0.38	-1.28***	0.26**	-1.56***	-0.44	-1.14***	0.05

Note: Coefficients on Post×WFH (1sd) and Post×GenAI (1sd) from the joint additive specification of Table 1, re-estimated 18 times — each time dropping all data for the focal 2-digit ONET-SOC major group identified by Code (first column) and Occupation Group Name (second column). Columns (1)–(4) report the Panel A specification on State×6-digit-Occupation×Month cells with State×Occupation and State×Month fixed effects, two-way clustered by 6-digit occupation and state. Columns (5)–(8) report the Panel B specification on the pre-aggregated 6-digit-shock-level BHJ panel from Table 1 restricted to shocks outside the focal 2-digit group, clustered by 2-digit occupation. Columns (1)–(2) and (5)–(6) use the junior share of new hires from Revelio Labs as the outcome; columns (3)–(4) and (7)–(8) use the share of Lightcast postings requiring at most three years of experience. Odd-numbered columns report the WFH coefficient and even-numbered columns the GenAI coefficient. Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.4: Alternative measures of WFH and GenAI exposure in the two-period diff-in-diff design

		Junior Worker Share of New Hires (Pp.)			Share of Job Postings Requiring ≤ 3 Years Experience (Pp.)		
WFH	GenAI	WFH (1)	GenAI (2)	WFH×AI (3)	WFH (4)	GenAI (5)	WFH×AI (6)
Panel A: State × Occupation × Month Cells							
Hansen et al.	Eloundou et al.	-1.63***	0.04	0.67***	-1.34***	0.35*	0.15
Hansen et al.	Felten et al.	-1.14***	-0.64	0.49	-1.69***	0.65**	0.36
Hansen et al.	Anthropic Economic Index	-1.78***	0.17	0.28	-1.17***	0.24	-0.07
Dingel–Neiman	Eloundou et al.	-2.54***	-1.17***	1.00*	-1.78***	0.01	-0.14
Dingel–Neiman	Felten et al.	-1.66**	-1.43***	0.43	-2.05***	0.04	0.13
Dingel–Neiman	Anthropic Economic Index	-3.72***	0.57	-0.95	-1.84***	0.22	-0.45
Panel B: Firm × 2-Digit Occupation × Year Cells							
Hansen et al.	Eloundou et al.	-1.57***	-0.51	-0.01	-1.10***	0.04	-0.05
Hansen et al.	Felten et al.	-1.36***	-0.88***	-0.00	-0.84***	-0.25	-0.30
Hansen et al.	Anthropic Economic Index	-1.82***	-0.13	-0.32	-1.07***	-0.06	0.00
Dingel–Neiman	Eloundou et al.	-1.22***	-1.23**	-0.15	-0.78***	-0.34**	-0.32
Dingel–Neiman	Felten et al.	-0.89**	-1.57***	-0.32	-0.83***	-0.25	-0.20
Dingel–Neiman	Anthropic Economic Index	-1.70***	-0.60*	-0.61*	-0.91***	-0.23***	-0.09

Note: Coefficients on Post×WFH (1sd), Post×GenAI (1sd), and Post×WFH×GenAI (1sd) from the joint specification with interaction (matching column (4)/(8) of Table 1), re-estimated under each of 2 WFH × 3 GenAI = 6 measure combinations. Hansen et al. denotes the WFH exposure from Hansen et al. (2023); Dingel–Neiman is the telework-ability index of Dingel and Neiman (2020). GenAI measures are the headline index of Eloundou et al. (2024), the index of Felten et al. (2023), and the Anthropic Economic Index of Appel et al. (2025). Columns (1)–(3) report the junior share of new hires and columns (4)–(6) the share of postings requiring ≤ 3 years experience. *Panel A* uses State×6-digit-Occupation×Month cells with State×Occupation and State×Month fixed effects, two-way clustered by 6-digit occupation and state. *Panel B* uses the pre-aggregated 6-digit-shock-level BHJ panel from Table 1, clustered by 2-digit occupation. Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.5: Residualized exposures preserve the WFH-dominant ranking

	Junior Worker Share of New Hires (Pp.)		Share of Job Postings Requiring ≤ 3 Years Experience (Pp.)	
	(1)	(2)	(3)	(4)
Panel A: State × Occupation × Month Cells				
Post × WFH (⊥ GenAI) (1sd)	-1.06***		-0.92***	
	(0.27)		(0.14)	
Post × GenAI (⊥ WFH) (1sd)		-0.29		0.22*
		(0.30)		(0.12)
Partial F-Statistic	15.5	0.9	44.6	3.4
Observations	2,312,186 cells / 120,659,812 hires		6,577,807 cells / 217,318,081 postings	
Panel B: Firm × 2-Digit Occupation × Year Cells				
Post × WFH (⊥ GenAI) (1sd)	-1.26***		-0.80***	
	(0.32)		(0.23)	
Post × GenAI (⊥ WFH) (1sd)		-0.35		0.04
		(0.63)		(0.18)
Partial F-Statistic	15.8	0.3	12.1	0.1
Observations	3,621,900 cells / 31,706,183 hires		2,065,266 cells / 58,940,327 postings	

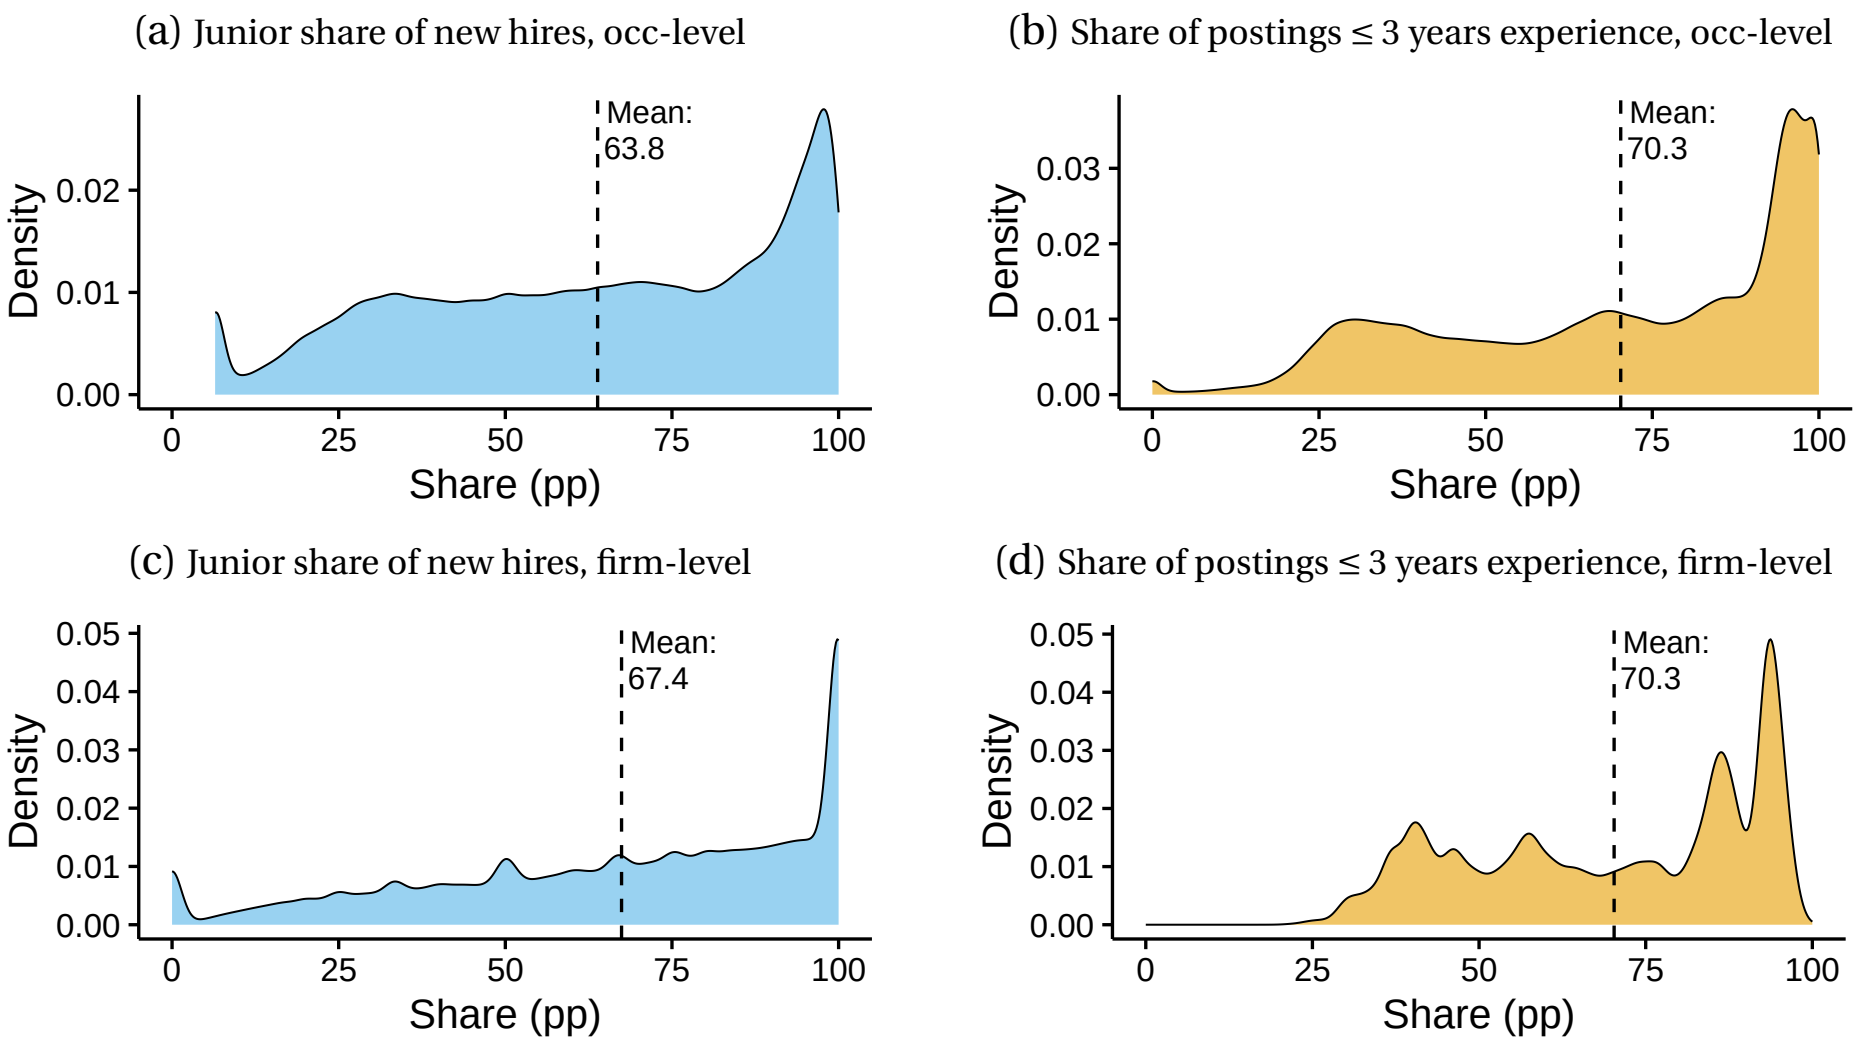
Note: Single-treatment DiD coefficients using the residualized exposures of the headline Hansen et al. (2023) + Eloundou et al. (2024) measure pair, constructed as in equation (E.4) of Section E.5. The residualized WFH series is the residual of the Hansen et al. WFH index after partialling out the Eloundou GenAI index at the occupation level, re-standardized within the analysis sample; the residualized GenAI series is symmetric. The odd-numbered columns regress on the post-interacted residualized WFH; the even-numbered columns regress on the post-interacted residualized GenAI. Columns (1)–(2) use the junior share of new hires; columns (3)–(4) the share of postings requiring three or fewer years of experience. *Panel A* uses State×6-digit-Occupation×Month cells, with State×Occupation and State×Month fixed effects, two-way clustered by occupation and state. *Panel B* uses Firm×2-digit-Occupation×Year cells with shift-share-robust inference following Borusyak et al. (2022), clustered by 2-digit occupation; residualization in Panel B is performed at the 6-digit shock level, weighted by the BHJ exposure weight. Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

Table E.6: Sensitivity to setting different Experience-requirement Thresholds in Job-postings

	≤ 1 yr	≤ 2 yr	≤ 3 yr	≤ 4 yr	≤ 5 yr	≤ 6 yr
Panel A: State × Occupation × Month cells (occ-level pre/post DiD)						
Post × WFH (1sd)	-1.76***	-1.01***	-1.07***	-0.84***	-0.94***	-0.83***
	(0.54)	(0.30)	(0.28)	(0.28)	(0.23)	(0.21)
Post × GenAI (1sd)	-0.07	-0.03	-0.01	0.18	0.10	0.05
	(0.34)	(0.24)	(0.20)	(0.18)	(0.13)	(0.12)
Observations	3,453,516	3,453,516	3,453,516	3,453,516	3,453,516	3,453,516
Panel B: Firm × 2-Digit Occupation (BHJ shock-level inference)						
WFH (1sd)	-1.14***	-1.43***	-1.12***	-1.06***	-1.09***	-1.01***
	(0.23)	(0.41)	(0.26)	(0.27)	(0.21)	(0.23)
GenAI (1sd)	0.20	0.34	-0.14	-0.13	-0.01	-0.02
	(0.26)	(0.34)	(0.22)	(0.24)	(0.16)	(0.16)
Shocks (6-digit ONET)	677	677	677	677	677	677

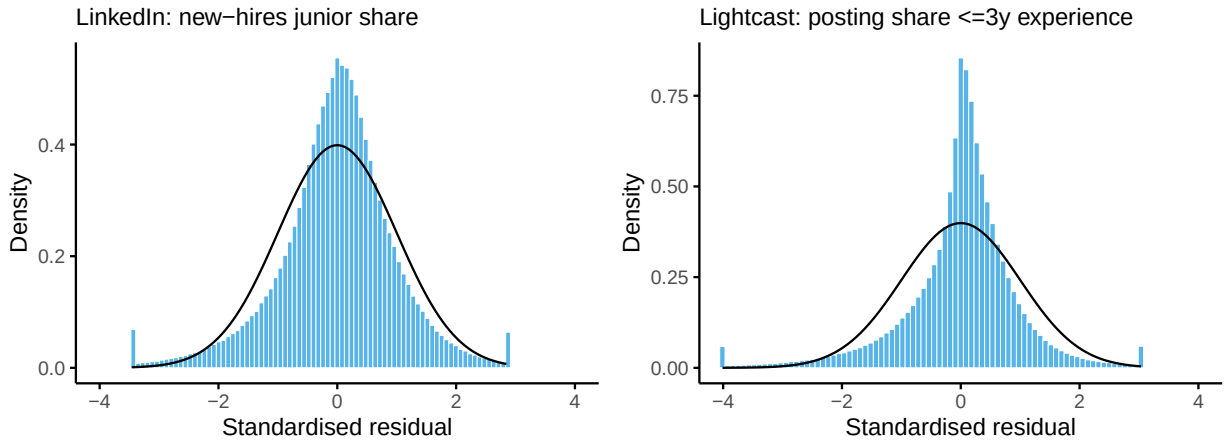
Notes: Re-estimates the joint job-postings specification from Table 1 under alternative experience-threshold definitions of "junior". The outcome in column k is the share of postings (with non-missing experience) requiring at most k years of experience. *Panel A* is the unit-level pre/post DiD at the State × Occupation × Month cell level, with State×Occupation and State×Month fixed effects, standard errors two-way clustered by ONET and State. *Panel B* reformulates the firm × 2-digit-occupation specification at the 6-digit-occupation shock level via Borusyak, Hull and Jaravel (2022) `ssaggregate`; long differences in the per-threshold share are aggregated to the shock level using pre-period firm-mix shares and country fixed effects partialled out, with standard errors clustered by 2-digit ONET. Exposures standardised at the panel level for both panels; winsorisation at 1/99 pct. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

Figure E.1: Kernel Densities of the Raw Share Outcomes used in our Main Empirical Specification (Table 1)



Note: Cell-weighted kernel densities of the two share outcomes used in Table 1, at the two cell aggregation levels used in that table. Panels (a) and (b): State×6-digit-Occupation×Month cells (Panel A of Table 1). Panels (c) and (d): Firm×2-digit-Occupation×Year cells (Panel B of Table 1). The dashed vertical line marks the cell-weighted mean; densities are weighted by the same total hires/postings used as the regression weight in Table 1. Lightcast cells are constructed from the IPW-corrected posting panel so that the share reflects the missingness correction described in Appendix C.2.

Figure E.2: Residuals from the Headline Joint OLS Specification Are Approximately Normal



Note: Histograms of standardised residuals from the headline joint difference-in-differences specification (Table 1, Panel A, column (3) for new hires and column (7) for postings), with a fitted $\mathcal{N}(0,1)$ density overlaid. Residuals are standardised by their weighted sample standard deviation. The horizontal axis is trimmed at ± 4 standard deviations for readability. *Left panel:* Revelio Labs new-hires junior share at the State×6-digit-Occupation×Month cell ($N = 1{,}009{,}585$), with sample skewness -0.45 and excess kurtosis 2.4. *Right panel:* Lightcast posting share requiring ≤ 3 years of experience at the same cell aggregation ($N = 2{,}002{,}405$), with sample skewness -0.88 and excess kurtosis 5.3 — mildly left-skewed and leptokurtic, broadly comparable to a $t(5)$. Both residual distributions are unimodal and centred at zero, consistent with the conditional-mean linearity assumption underlying our headline OLS-on-percentage-points specification (see Appendix E.6 for further discussion).

Table E.7: Robustness: Alternative Specifications (Joint WFH + GenAI)

	Occ-Level (State × 6-Digit Occ × Month)		Firm-Level (Firm × 2-Digit Occ × Year)	
	Junior share of new hires	Share of postings ≤ 3 years exp.	Junior share of new hires	Share of postings ≤ 3 years exp.
	(1)	(2)	(3)	(4)
Panel A: Unweighted OLS (pp)				
Post × WFH (1sd, pp)	-1.41***	-0.97***	-2.26***	-1.04***
	(0.30)	(0.16)	(0.65)	(0.17)
Post × GenAI (1sd, pp)	-0.52**	-0.03	-0.20	0.25**
	(0.20)	(0.11)	(0.64)	(0.11)
Partial F-Statistic	21.5/6.4	39.0/0.1	12.1/0.1	37.9/5.5
Panel B: Poisson PML (Papke–Wooldridge), Weighted by Cell Size				
Post × WFH (1sd, pp)	-2.98***	-1.55***	-4.20***	-1.74***
	(0.66)	(0.23)	(1.12)	(0.22)
Post × GenAI (1sd, pp)	0.52	0.37***	0.96	0.48***
	(0.57)	(0.12)	(1.01)	(0.15)
Partial F-Statistic	19.6/0.8	42.8/9.1	13.1/0.9	58.9/10.2
Panel C: OLS on Log-Odds Transform, $\log\big(s/(1-s)\big)$				
Post × WFH (1sd, pp)	-0.28	-0.80***	-2.71***	-0.89***
	(0.45)	(0.28)	(0.66)	(0.28)
Post × GenAI (1sd, pp)	-0.30	0.64***	-0.39	0.39
	(0.50)	(0.21)	(0.62)	(0.25)
Partial F-Statistic	0.4/0.4	7.9/9.8	16.9/0.4	10.3/2.5

Note: Robustness check for the joint WFH + GenAI specification of Table 1 (column 3 of each block). *Panel A* is unweighted OLS on the outcome in pp; *Panel B* is the Papke–Wooldridge fractional Poisson PML on the same outcome, weighted by cell size; *Panel C* is OLS on the log-odds transform of the outcome. Cols. (1)–(2) match Table 1 Panel A's State×6-digit-Occupation×Month cells; cols. (3)–(4) match Panel B's Firm×2-Digit-Occupation×Year cells (at the unit rather than BHJ shock level). All coefficients are reported as the implied marginal effect in pp at the cell-weighted mean of the outcome, with delta-method standard errors; conversion formulas, 1/99 outcome winsorisation, and the rationale for the fractional-Poisson formulation are in Appendix E.6. Significance stars are based on the delta-method z-statistic. Partial F is the squared t-ratio of the native-scale focal coefficient (WFH/GenAI). Significance codes: *** $p < 0.01$, ** $p < 0.05$, * $p < 0.1$.

F Additional Robustness Exercises

F.1 Bin-Depth Coarsening Diagnostic for the Co-Treatment

Figure F.1 reports the pre–post difference-in-differences coefficient on each focal exposure as the co-treatment is absorbed at progressively finer resolutions, from no control through 10 quantile bins. The WFH coefficient (left column) is essentially unchanged across resolutions in both data sources. The GenAI coefficient (right column) attenuates monotonically toward zero in both data sources as the WFH control is absorbed at finer resolutions.

F.2 Computing the Variance-Inflation and Condition-Number Diagnostics

This subsection records the linear-algebra detail behind the VIF and CN diagnostics referenced in Section 5 and reported alongside the joint specifications in Section 6. For a focal treatment column x_h, the variance inflation factor is

$$\mathrm{VIF}_h = \frac{1}{1 - \mathscr{R}_h^2}, \tag{F.1}$$

where $\mathscr{R}_h^2$ is from regressing x_h on the other focal treatment columns. The condition number is the system-wide analogue,

$$\mathrm{CN} = \sqrt{\frac{\lambda_{\max}(R_{\mathscr{F}})}{\lambda_{\min}(R_{\mathscr{F}})}}, \tag{F.2}$$

where $R_{\mathscr{F}}$ is the correlation matrix of the focal treatment block. In specifications with two focal regressors and residualized correlation ρ, these diagnostics reduce to

$$\mathrm{VIF} = \frac{1}{1 - \rho^2}, \tag{F.3}$$

$$\mathrm{CN} = \sqrt{\frac{1 + |\rho|}{1 - |\rho|}}. \tag{F.4}$$

In practice, we compute these diagnostics from the residualized, weighted treatment columns that identify the reported coefficients, not from the raw exposure indices. Let $X_{\mathscr{F}}$ collect the focal treatment columns, let A collect fixed effects and non-focal controls, and let W be the

regression-weight matrix. We form

$$\widetilde{X}_{\mathscr{F}} = \left[I - A(A'WA)^{-1}A'W \right] X_{\mathscr{F}}, \tag{F.5}$$

$$H_{\mathscr{F}} = \widetilde{X}'_{\mathscr{F}} W \widetilde{X}_{\mathscr{F}}, \tag{F.6}$$

$$R_{\mathscr{F}} = S^{-1/2} H_{\mathscr{F}} S^{-1/2}, \qquad S = \operatorname{diag}(H_{\mathscr{F}}). \tag{F.7}$$

We report

$$\mathrm{VIF}_h = \left(R_{\mathscr{F}}^{-1}\right)_{hh}, \tag{F.8}$$

$$\mathrm{CN} = \sqrt{\frac{\lambda_{\max}(R_{\mathscr{F}})}{\lambda_{\min}(R_{\mathscr{F}})}}. \tag{F.9}$$

Thus, VIF and CN measure collinearity in the identifying variation that remains after fixed effects, weights, controls, matching weights where relevant, and aggregation.

F.3 Selection-on-Unobservables Sensitivity

This subsection asks how robust the single-treatment WFH and GenAI difference-in-differences coefficients of Table 1 are to omitted-variable bias. This is in the spirit of Altonji et al. (2008) and Oster (2019), but given the high-dimensional fixed effects we follow the approach of Cinelli and Hazlett (2020).

Cinelli and Hazlett (2020) parameterise an unobserved confounder W directly through its partial R^2 with the focal exposure Post $\times Z^k$ and with the outcome, both residualized on the fixed effects of (9) and the co-treatment Post $\times Z^\ell$, and report the robustness value

$$RV_{q=1} = \frac{1}{2}\left(\sqrt{f_\beta^4 + 4f_\beta^2} - f_\beta^2\right), \qquad f_\beta \equiv \frac{|\widehat{\beta}|}{\mathrm{SE}(\widehat{\beta})\sqrt{\mathrm{df}}}, \tag{F.10}$$

where df are the residual degrees of freedom. $RV_{q=1} \in [0,1]$ is the smallest partial R^2 that W would need to share with both treatment and outcome to drive $\widehat{\beta}$ to zero; higher values imply a greater robustness. A value of 0.10, for instance, says that any unobservable accounting for at least 10% of the residual variation in both treatment and outcome would suffice.

Table F.1 reports $RV_{q=1}$ for the single-treatment WFH (column 1) and GenAI (column 2) designs, with the other exposure entering as a candidate omitted variable. In the occupation-level exposure designs of Table 1, the WFH robustness value is 3–5 times larger than the GenAI one: 0.05 vs. 0.01 in Panel A (postings outcome) and 0.09 vs. 0.03 in Panel B (junior-hire share). A

confounder of 1–3% partial-R^2 in both treatment and outcome would suffice to eliminate the GenAI effect, while the WFH effect would require a 5–9% confounder. Panels C and D apply the same diagnostic to the firm-level actual-WFH design of Table 2, for which there is no GenAI counterpart.

F.4 Simulating Classical Measurement Error

Another potential concern is that the measurement quality of our preferred WFH and GenAI exposure measures may be asymmetric. In this subsection we run a simulation to test whether classical measurement error in the GenAI index could explain our empirical results. In each replication we set the true WFH effect to zero, $\beta^{\star}_{\text{WFH}} = 0$, and assign all explanatory power to GenAI. Specifically, we construct a noisy GenAI measure:

$$Z^{\text{obs}}_{\text{GenAI}} = Z_{\text{GenAI}} + \varepsilon_o, \qquad \varepsilon_o \sim \mathcal{N}(0, \nu^2), \tag{F.11}$$

where ε_o is drawn once per ONET occupation and the resulting series is re-standardised to unit variance. The parameter $\nu = \sigma_\varepsilon / \sigma_Z$ is the noise-to-signal standard-deviation ratio; the reliability of the index is $R = 1/(1+\nu^2)$. We then consider how often the joint regression nonetheless return a statistically significant WFH coefficient—i.e. how often does the noise in GenAI cause WFH to be falsely credited with the effects? We consider two calibrations of the true GenAI effect $\beta^{\star}_{\text{GenAI}}$. **Panel A** (LinkedIn hires) sets $\beta^{\star}_{\text{GenAI}}$ equal to the sum of the two empirical coefficients, $-(|\hat{\beta}^{\text{emp}}_{\text{WFH}}| + |\hat{\beta}^{\text{emp}}_{\text{GenAI}}|)$, so that GenAI alone accounts for the full observed decline. **Panel B** (Lightcast postings) sets $\beta^{\star}_{\text{GenAI}}$ to the single-treatment GenAI coefficient on postings—the magnitude one recovers before WFH is added as a control.

The results are reported in Table F.2. In Panel A, the false-positive rate for WFH is essentially zero whenever the GenAI index is at least 64% reliable ($R \geq 0.64$), and exceeds the nominal 5% level only when noise variance equals signal variance ($\nu = 1$, $R = 0.50$). In Panel B, the false-positive rate is 0% at every reliability level tested: even at $R = 0.50$ the median simulated WFH coefficient is only -0.20 pp and is never significant, far below the empirical estimate of -1.10 pp. These findings suggest to us that classical attenuation bias in the GenAI index would need to be implausibly large to generate the WFH-dominant pattern we observe if the true underlying DGP were only driven by GenAI exposure.

F.5 Obviously Related IV to Address Measurement Error

Our next exercise seeks to address potential measurement error in the GenAI exposure side by using the obviously related IV tools proposed in Gillen et al. (2019). For each of our three measures of GenAI exposure, we take one focal measure and use as instruments the two other measures. Table F.3 reports the OLS and ORIV coefficients. The metric of Gillen et al. (2019) used to assess the presence of measurement error is the implied reliability: $\hat{R} = \hat{\beta}_{\text{AI}}^{\text{OLS}} / \hat{\beta}_{\text{AI}}^{\text{ORIV}}$, which we report alongside the reduced form and IV results. Under classical measurement error $\hat{R} \in (0,1)$ and equals the share of the observed-measure variance attributable to the common signal after partialling out the co-treatment of WFH exposure and the fixed effects.

The implied reliability $\hat{R}$ reported in Table F.3 falls outside the classical-error range $(0,1)$ in every row, suggesting a classical-measurement-error story is again unlikely to drive our finding that WFH exposure is a driver of the fall in the relative hiring intensity of junior-workers. In terms of the actual regression coefficients, the only instance where WFH exposure does not have a larger and significant effect on the junior-share of new hires is when we instrument the Anthropic Economic Index. We do not put much weight on this finding, given the very weak reduced form relationship and the fact that it arises only in Panel A, and only for this specific combination. We do believe that the ORIV approach in general holds great promise in the context of many related exposure measures, and should be more widely deployed.

Table F.1: Selection-on-Unobservables Diagnostics for Single-Treatment Designs

	Single-treatment WFH design (1)	Single-treatment GenAI design (2)
Panel A: Share of Job Postings Requiring ≤ 3 Years Experience		
Exposure Design at the Occupation×State×Month level (see Table 1 Panel A)		
Cinelli–Hazlett RV ($q = 1$, $\alpha = 1$)	0.05	0.01
Cinelli–Hazlett RV ($q = 1$, $\alpha = 0.05$)	0.05	0.01
Panel B: Junior Worker Share of New Hires		
Exposure Design at the Occupation×State×Month level (see Table 1 Panel A)		
Cinelli–Hazlett RV ($q = 1$, $\alpha = 1$)	0.09	0.03
Cinelli–Hazlett RV ($q = 1$, $\alpha = 0.05$)	0.09	0.02
Panel C: Share of Job Postings Requiring ≤ 3 Years Experience		
Actual WFH Design at the Firm×Occupation×Year level (see Table 2 Panel A)		
Cinelli–Hazlett RV ($q = 1$, $\alpha = 1$)	0.01	—
Cinelli–Hazlett RV ($q = 1$, $\alpha = 0.05$)	0.01	—
Panel D: Junior Worker Share of New Hires		
Actual WFH Design at the Firm×Occupation×Year level (see Table 2 Panel A)		
Cinelli–Hazlett RV ($q = 1$, $\alpha = 1$)	0.02	—
Cinelli–Hazlett RV ($q = 1$, $\alpha = 0.05$)	0.01	—

Notes: The table reports the Cinelli and Hazlett (2020) robustness value $RV_{q=1}$ for each single-treatment difference-in-differences design. The takeaway is that the WFH-design $RV_{q=1}$ is 3–5× larger than the GenAI-design $RV_{q=1}$ in the exposure designs (Panels A and B). Panels A and B correspond to the occupation-level exposure designs of Table 1 (Panel A); Panels C and D correspond to the firm-level realised-WFH design of Table 2 (Panel A). There is no firm-level GenAI analogue, so only the WFH column is populated in Panels C and D. Oster (2019) δ^* is discussed but not reported numerically in Appendix F.3.

Table F.2: Misattributing a true GenAI effect to WFH requires implausibly noisy GenAI exposure (Monte Carlo simulations)

Severity	Noise in GenAI index		Median estimate (pp)		Misattribution
	Noise/Signal SD ratio ν	Reliability R	$\hat{\beta}_{\text{WFH}}$	$\hat{\beta}_{\text{GenAI}}$	rate (%)
	(1)	(2)	(3)	(4)	(5)
Panel A: Junior Worker Share of New Hires (pp). Simulated data-generating process: WFH effect = 0; GenAI effect = −2.07 pp.					
None (clean signal)	0.00	1.00	0.00	−2.07	0.0
Slight	0.25	0.94	−0.18	−1.87	0.0
Modest	0.50	0.80	−0.52	−1.52	0.0
Moderate	0.75	0.64	−0.82	−1.17	1.6
Equal noise & signal	1.00	0.50	−1.02	−0.93	7.0
Panel B: Share of Job Postings Requiring ≤ 3 Years Experience (pp). Simulated data-generating process: WFH effect = 0; GenAI effect = −0.50 pp (single-treatment GenAI coefficient on postings).					
None (clean signal)	0.00	1.00	0.00	−0.51	0.0
Slight	0.25	0.94	−0.04	−0.46	0.0
Modest	0.50	0.80	−0.11	−0.39	0.0
Moderate	0.75	0.64	−0.18	−0.31	0.0
Equal noise & signal	1.00	0.50	−0.20	−0.27	0.0

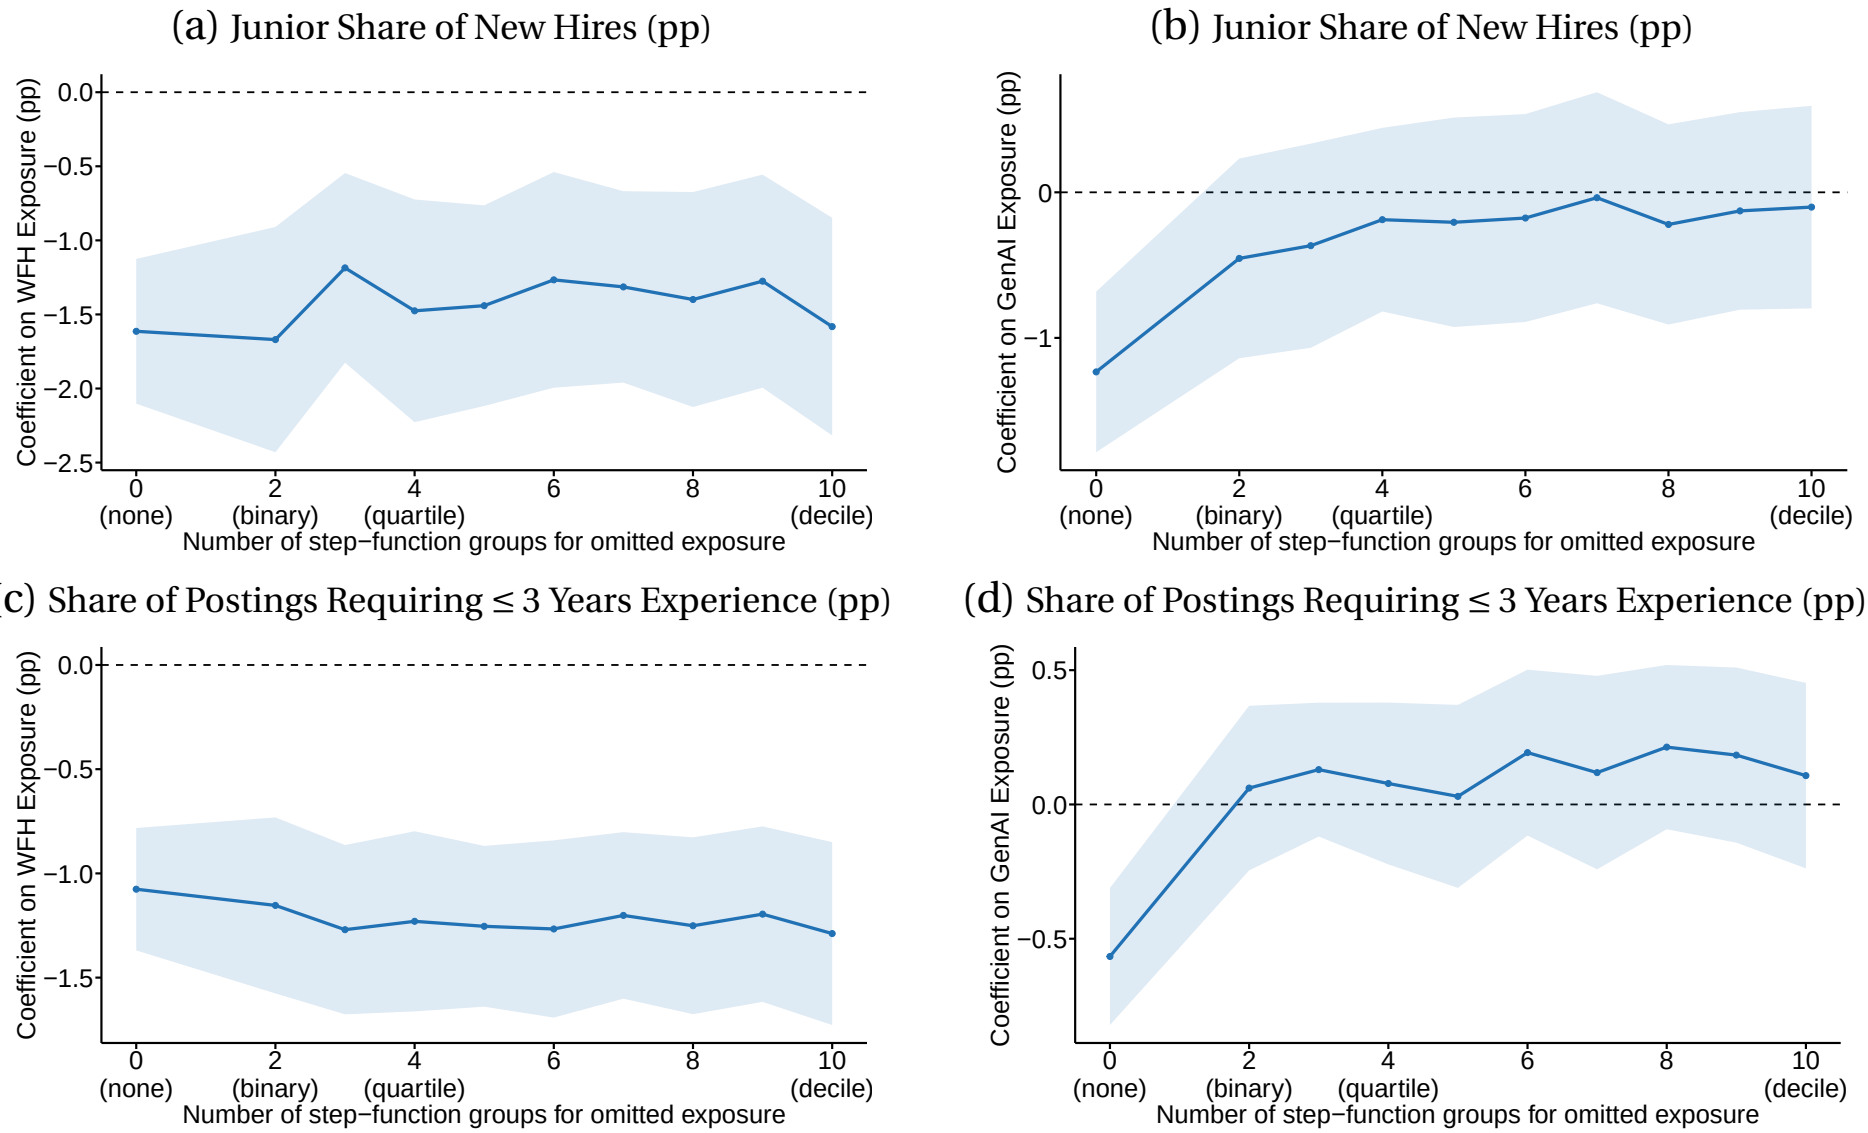
Notes: Each row reports a Monte Carlo simulation asking how noisy the GenAI index must be before the joint regression incorrectly assigns a true negative GenAI effect to WFH. We simulate the outcome with a true WFH effect of zero, add Gaussian noise to the observed GenAI index, re-standardise it to unit variance, and re-estimate the joint specification 500 times; the WFH index is held fixed throughout. Panel A loads the full empirical magnitude onto a negative GenAI effect ($\beta^*_{\text{GenAI}} = -(|\hat{\beta}^{\text{emp}}_{\text{WFH}}| + |\hat{\beta}^{\text{emp}}_{\text{GenAI}}|)$); Panel B sets β^*_{GenAI} equal to the (negative) single-treatment GenAI coefficient on postings, the natural "GenAI-is-the-only-driver" alternative. The misattribution rate is the share of replications in which $|\hat{\beta}_{\text{WFH}}| \geq 3\,|\hat{\beta}_{\text{GenAI}}|$, WFH is significant at 5%, and GenAI is not. Panel A's misattribution rate stays below 2% until $R \leq 0.64$ and becomes appreciable only once noise variance equals signal variance ($\nu = 1$, $R = 0.5$). Panel B's misattribution rate is 0% across all reliability levels considered: even at $R = 0.50$ the median simulated WFH coefficient is only −0.20 pp and is not statistically significant, far smaller in magnitude than the empirical −1.10 pp. $\nu = \sigma_\varepsilon / \sigma_{Z_{\text{GenAI}}}$ is the noise-to-signal SD ratio (distinct from the learning-rate λ in Section 3) and $R = 1/(1+\nu^2)$ is the reliability ratio. Calibration uses the empirical occupation × state × month panel ($\rho_{\text{WFH,GenAI}} \approx 0.67$ in the standardised exposure).

Table F.3: Obviously-Related IV (ORIV) joint specification across three GenAI exposure measures

	Reduced Form OLS				Obviously Related IV				IV diagnostics	
Measure	$\hat{\beta}_{\text{AI}}$	(SE)	$\hat{\beta}_{\text{WFH}}$	(SE)	$\hat{\beta}_{\text{AI}}$	(SE)	$\hat{\beta}_{\text{WFH}}$	(SE)	First-stage F	Implied $\hat{R}$
(1)	(2)	(3)	(4)	(5)	(6)	(7)	(8)	(9)	(10)	(11)
Panel A: Junior Worker Share of New Hires										
Eloundou et al.	−0.52	(0.39)	-1.57^{***}	(0.39)	−0.35	(0.76)	-1.69^{***}	(0.61)	407,345	1.50
Felten et al.	-0.73^{**}	(0.33)	-1.43^{***}	(0.37)	−0.15	(1.20)	-1.84^{**}	(0.89)	187,301	4.99
Anthropic Economic Index	0.31	(0.39)	-2.13^{***}	(0.33)	-1.98^{*}	(1.10)	−0.76	(0.78)	123,163	−0.16
Panel B: Share of Job Postings Requiring ≤ 3 Years Experience										
Eloundou et al.	−0.05	(0.28)	-1.60^{***}	(0.39)	0.74	(0.59)	-2.19^{***}	(0.78)	574,775	−0.07
Felten et al.	0.83^{**}	(0.37)	-2.25^{***}	(0.56)	−0.07	(0.52)	-1.58^{**}	(0.62)	418,313	−11.07
Anthropic Economic Index	−0.02	(0.27)	-1.63^{***}	(0.51)	0.85	(0.56)	-2.18^{***}	(0.59)	231,026	−0.02

Notes: Each row uses one of three standardized GenAI exposure measures m_k. Columns (2)–(5) report the OLS joint-specification coefficients on Post × m_{ko} and Post × Z_o^{WFH} from the same regression. Columns (6)–(9) report the AI and WFH coefficients from a separate ORIV regression per row, in which the post-interacted AI exposure is instrumented by the post-interacted mean of the other two standardized AI measures, following Gillen et al. (2019). Column (10) reports the first-stage F. Column (11) reports the implied reliability $\hat{R} = \hat{\beta}_{\text{AI}}^{\text{OLS}} / \hat{\beta}_{\text{AI}}^{\text{ORIV}}$; under classical measurement error $\hat{R} \in (0, 1)$ and equals the share of the observed-measure variance attributable to the common signal (the same reliability quantity R used in the Monte Carlo simulation of Appendix F.4). Missing occupations in any source measure are imputed by hierarchical 5/3/2-digit ONET means so the analysis covers the headline occupation sample. Panel A outcome is the share of new hires classified as entry-level or junior, in percentage points, from the new-hires data; Panel B outcome is the share of job postings requiring three years of experience or fewer, in percentage points, from the job-postings data. All specifications include market×occupation and market×month fixed effects, sample weights (new-hire weights in Panel A, posting weights in Panel B), and two-way clustered standard errors at the 7-digit ONET and market level. Significance codes: ***: 0.01, **: 0.05, *: 0.1.

Figure F.1: Robustness to Non-Linear Co-Treatment Controls: Bin-Depth Diagnostic (Occupation×State×Month)



Note: Each panel plots the pre/post DiD coefficient on the focal exposure as the co-treatment is absorbed at different step-function resolutions. The x-axis is the number of equal-sized quantile bins used to absorb the co-treatment (0 = no control; 10 = deciles). Left column: WFH focal, GenAI absorbed. Right column: GenAI focal, WFH absorbed. Top row (panels a, b): Revelio Labs new-hires junior share. Bottom row (panels c, d): Lightcast posting share requiring ≤ 3 years experience. The panel uses the 6-digit O*NET-SOC×state×calendar-month design from Table 1 Panel A, with region×occupation and region×month fixed effects, clustered by 6-digit occupation and region. Shaded regions give 95% confidence intervals.