← Back to Writing

WORKING PAPER, PARTS I TO III

The Last Human Task

What 2.53 billion job postings say about the broken rung, July 2019 to May 2026. A working paper on AI and junior hiring, with interactive charts.

By Adi Gupta · Mon, Nov 9, 2026

Economics of AI | Labor markets | The broken rung | Econometrics | Original research | Working paper | Job postings | Revelio Labs via WRDS

WORKING PAPER, PARTS I TO III · THIS VERSION: NOVEMBER 9, 2026

Aditya (Adi) Gupta

adi@adya.one

Wharton MBA for Executives, 2028

Data: Revelio Labs via WRDS

The claim under test is simple and loud: AI does the junior work first, so firms stop hiring juniors, and the bottom rung of white-collar ladders comes off. Stanford's payroll study found the first-stage symptom in 2023 to 2025 employment. This paper asks whether the symptom is visible in hiring intent, in every job posting Revelio Labs captured worldwide over the two years in which agentic models arrived, with a

esign that cannot mistake a sector slump, a seasonal cycle, or a change in data coverage for a broken rung.

In short, it is not visible yet. Exposed occupations posted more jobs, not fewer; their junior-titled postings did not fall behind their senior-titled postings; their low-pay postings grew relative to their high-pay postings; and the pay offered in them drifted down relative to other occupations. Layoffs tilted toward exposed firms in the 2023 wave and tilted back after the agentic wave. If AI is pressing on these occupations in 2024 to 2026, it shows up first in the price of a posting, not in the count of junior postings. Every number below carries a confidence tag, and the limitations section says what this cannot show.

Parts II and III (added November 2026) push the window back to 2019 and add people to postings. With 2.53 billion postings from July 2019, exposed occupations' LinkedIn-sourced postings fell by more than half relative to unexposed ones through 2023 to 2025 and recovered in 2026, but the decline is a computer-occupations story that began before ChatGPT, within occupation groups it is a quarter with an interval reaching zero, two of three seniority measures show no tilt against juniors, while Revelio's own seniority field, which now runs to May 2026, shows a within-occupation tilt of about -16 points in exposed computer and IT occupations from 2024 on that the title measure does not confirm; offered pay fell. In the sources that reach back to 2019, the exposed-occupation gradient moved by ChatGPT-sized amounts between 2019 and 2022 for reasons that had nothing to do with AI (the COVID shock and coverage drift), which is the scale against which the 2023 decline has to be read. In Revelio's firm-level headcount for 5,348 public firms, AI-exposed firms became *more* junior-heavy after ChatGPT, by shedding senior staff, and the one broken-rung signal in

ıe data, a smaller junior share of new hires at exposed firms, lasted about eighteen months from late 2022 and had reversed by 2026.

Main findings

Each row is one design: what was held constant, what the data showed, and how I read it. "Exposed" means a high score on the Eloundou et al. rating of how much of an occupation's work a language model could do; "junior" and "senior" are defined in Section 3.

> Power, stated once: the title-tier estimate has a standard error of about 0.28 and the pay-band estimate about 0.05, on a scale where the Stanford payroll result would be roughly −0.05 per year. The pay-band test can see an effect that size; the title test cannot. Read the headline as "no sign of a broken rung in hiring intent, and some sign of the opposite", not as proof that none exists.

---

1. Why this question, now

Three things happened between mid-2024 and mid-2026. Frontier models went from chat to agents: computer use (Oct 2024), reasoning models (o1, R1), and coding agents that ship software (Claude Code, Feb to May 2025), followed by GPT-5, Sonnet 4.5, Opus 4.5 and 4.6, GPT-5.4 and the Claude 5 generation. Second, the labor-economics literature produced its first credible displacement evidence: Brynjolfsson, Chandar and Chen (2025) used ADP payroll records and found a roughly 13% relative employment decline for 22 to 25 year olds in the most AI-exposed occupations, with older workers

naffected high: published paper. Third, the public argument escalated. Anthropic's CEO said in May 2025 that AI could eliminate half of entry-level white-collar jobs within one to five years high: on-record interview, Axios, May 2025; MIT's Iceberg Index put the technically exposed share of US wage value at about 11.7% medium: press coverage of the Nov 2025 release.

Payroll data shows who was employed. Postings show who firms were trying to hire, months earlier and at occupation resolution. If the broken rung is a hiring decision, it should appear in postings first. That is the test.

---

2. Data

Revelio Labs job postings via Wharton's WRDS subscription: two extracts, June 2024 to June 2025 (72.8 GB, 483.9M rows) and June 2025 to June 2026 (122.5 GB, 837.6M rows), 1.32 billion postings after removing the 2.65M duplicates where the extracts overlap. Every row carries an O*NET-SOC occupation code, a normalized role title (Revelio's 17,000-category taxonomy), country, firm identifier, predicted salary, stated salary where posted, expected hires, and source flags (LinkedIn vs company sites and aggregators). The slim schema has no seniority field, no metro, and no posting text. US postings: 356.5M. Also used: 62,601 WARN-style layoff tranches 1990 to 2026 with firm identifiers, and the Anthropic Economic Index, Eloundou et al. and Felten exposure files (public).

Processing: one streaming pass per extract with DuckDB into columnar parquet (about 10 minutes on an 18-core laptop for 195 GB), then monthly cells by country, occupation, pay band, title tier and source. The histogram of predicted pay per occupation-month is stored, so

ny pay-band definition can be recomputed without re-reading the raw files. Code and aggregates are described in the reproducibility note.

---

3. Measuring seniority without a seniority field

Two constructions, both reported.

- **Fixed real pay bands (primary).** Revelio's predicted salary is present on 99% of US postings and correlates 0.95 in logs with stated salary where both exist. For each occupation and source, the 33rd and 67th percentiles of predicted pay over July to December 2024 define a low band and a high band; the thresholds grow at 3.5% a year so nominal wage growth does not move postings across bands. "Junior" is a posting in the low band, "senior" one in the high band. Validation on postings whose titles state seniority: in software development 94% of intern/junior-titled postings fall in the bottom tercile and 68% of senior/lead-titled postings in the top; for lawyers 82% and 50%; for financial analysts 82% and 11% (analyst ladders are flat). The band is a pay proxy for rung, and it works where the ladder is steep.
- **Title keywords (secondary).** Intern, entry, junior, trainee, apprentice, graduate vs senior, lead, principal, staff, manager, director, head, VP, chief. Covers about 15% of US postings, is dominated by a few very large categories (one tax-traineeship family alone), and is strongly seasonal. Reported for completeness.

A third candidate, within-year pay terciles, was tried and discarded: it drifts with wage growth inside each year and resets at year boundaries, which produced sign flips in release event studies that disappear under the fixed-band definition.

4. Design

The unit is an occupation (O*NET), a tier (junior or senior), a posting source, and a month; the outcome is the log of postings in the cell. Exposure is the Eloundou et al. (2024) GPT-4-rated beta score (0 to 1), with the Anthropic Economic Index observed exposure and the Felten AIOE index as alternatives. The estimating equation is a triple difference:

$$\ln n_{o,s,\mathrm{src},t} = \alpha_{o,s,\mathrm{src}} + \alpha_{o,t,\mathrm{src}} + \alpha_{s,t,\mathrm{src}} + \beta \cdot \mathrm{Exposure}_o \times \mathrm{Junior}_s \times \mathrm{Years}_t + \varepsilon_{o,s,\mathrm{src},t}$$

The occupation-month-source effect absorbs anything that hits an occupation in a month, juniors and seniors alike: a sector slump, a rate shock, a change in how many postings Revelio captured for that occupation from that source. The tier-month-source effect absorbs anything that hits all juniors at once, including internship seasons and an economy-wide junior freeze. β is identified only by junior postings moving relative to senior postings in the same occupation, more so in exposed occupations, over time. A year-over-year variant adds occupation-tier-calendar-month effects so that only within-calendar-month changes between 2024 to 25 and 2025 to 26 identify β. Standard errors are clustered by occupation (800 to 900 clusters). Post-period dummies (March 2025, the agentic coding wave; August 2025, GPT-5) and event-time fans around each major release are estimated on the same panel.

> Reading β: it is the extra annual log change in junior postings relative to senior postings for a fully exposed occupation (exposure 1) compared with an unexposed one (exposure 0). A broken rung is β well below zero.

---

5. Results

5.1 Occupation level

Title tiers. On the 145 occupations with enough junior- and senior-titled postings in every month, β = β = -0.42 (s.e. 0.28, p = 0.14). Weighted by baseline postings it is -0.03 (s.e. 0.53); with the Anthropic Economic Index as the exposure measure +0.30 (s.e. 0.27); with the Felten AIOE index +0.013 per s.d. (s.e. 0.045). Nothing here is distinguishable from zero and the signs disagree across measures. Comparing only the most and least exposed quintiles, junior-titled postings in the most exposed occupations grew faster after March 2025 (+0.31 log points, s.e. 0.08, 72 occupations).

Pay-band tiers. On 894 occupations:

The pay-band β is positive in every source and specification, significant in LinkedIn-sourced postings and in the year-over-year version, and not significant once occupations are weighted by size. Two post-period dummies agree: exposure × junior × post-March-2025 is +0.08 (s.e. 0.05), post-August-2025 is +0.18 (s.e. 0.05). A positive β means low-pay postings in exposed occupations grew relative to high-pay postings. Some of that is pay rather than people: Section 5.7 shows offered salaries in exposed occupations fell relative to other occupations, and a posting whose offered pay drops crosses from the high band

ıto the low band without anyone hiring a junior. The sensitivity check makes the point: with the band thresholds held flat instead of grown at 3.5% a year, β rises to +0.26 (s.e. 0.05), which is what pay drift in exposed occupations would do. The title-tier result, which cannot be moved by pay, is the cleaner headcount test, and it is null.

5.2 Release dates: no breaks

For each major release, the same regression with month-relative-to-release dummies replacing the trend (window of six months either side, month −1 omitted). The fans are jagged in both directions, pre and post, because occupation-specific seasonality is not removed by the tier-month effect. No release shows a clean step. The pooled post coefficients run from −0.47 to +0.36, changing sign with the season in which the release fell.

5.3 Levels: exposed occupations posted more

Dropping the tier dimension and regressing log total postings on exposure × years with occupation-source and month-source effects gives $\beta = +0.32$ (s.e. 0.05, $p < 0.001$) pooled, +0.22 (s.e. 0.06) on LinkedIn-sourced postings and +0.42 (s.e. 0.06) on other sources, across 917 occupations. Exposed occupations gained postings relative to unexposed ones over the window, in LinkedIn-sourced and other postings alike. This is the expansion check in the research design (H5): cheaper cognition can expand demand for the occupations that use it. It is also consistent with a cyclical rebound in white-collar hiring from the 2023 to 2024 trough, which the data cannot separate without a pre-2024 baseline.

5.4 Firms

Each firm's exposure is the posting-weighted mean Eloundou score of its own US postings (15,289 firms with at least 50 postings). Regressing the low-pay share of a firm's monthly postings on exposure × years with firm and month effects gives β = −0.045 (s.e. 0.008, p < 0.001; 238,825 firm-months, 15,092 firms): a firm at exposure 0.8 versus one at 0.4 moved its posting mix away from low-pay roles about 1.8 percentage points a year faster. With title tiers the coefficient is −0.003 (s.e. 0.005). The two results are not in conflict with 5.1: occupation-level composition is stable while within-firm composition shifts, which is what you would expect if exposed firms are re-weighting toward higher-pay roles across occupations rather than cutting junior rungs within them.

Revelio's workforce-dynamics table (firm × month × seniority 1 to 7) would carry a direct employment test. The extract on disk covers two months (April and May 2026), so it is a cross-section only; the Part II note below asks for the full series.

5.5 Layoffs

Layoff tranches from 1990 to 2026 join to the same firm exposure through the firm identifier (5,698 firms with any layoff since 2019). A linear probability model of any layoff event in a firm-month on exposure × post-ChatGPT (Dec 2022) and exposure × post-agents (Mar 2025), with firm and month effects: exposure × post-ChatGPT +0.056 (s.e. 0.005, p < 0.001), exposure × post-agents −0.018 (s.e. 0.006, p = 0.001); 512,550 firm-months. Exposed firms laid off more in 2023 to 2024 than their pre-2023 pattern implied, and that tilt partly reversed after March 2025. Exposure here is measured from 2024 to 2026 postings, and exposed firms are disproportionately tech and finance,

o this is as much a story about the rate-driven 2023 wave as about AI.

5.6 Thirty countries: no wage gradient

The same triple difference per country (within-year pay-tercile tiers), then the country coefficients against the country's mean predicted salary. If AI displaces juniors where humans are expensive first, β should fall as wages rise. It does not: the weighted slope of β on log wage level is +0.040. Of 30 countries, 5 have a negative point estimate, 3 significantly so (Italy, France, Mexico), and 9 are significantly positive. The US coefficient is +0.05 (s.e. 0.10).

5.7 Wage offers

On LinkedIn-sourced postings with a stated salary range (about a quarter of them), the log of the mean stated salary in an occupation-tier-month cell, same triple difference: β = −0.093 (s.e. 0.038, p = 0.016) on the junior-versus-senior contrast and −0.049 (s.e. 0.022, p = 0.025) on the level. Junior wage offers in exposed occupations slipped relative to senior offers, and exposed occupations' offers slipped slightly overall. Selection into stating a salary is not random, which keeps this at low-to-medium confidence.

5.8 New roles

Revelio maps raw titles to a fixed taxonomy, so genuinely new titles cannot appear; growth within AI-labelled categories can. AI- or ML-titled roles (AI, machine learning, LLM, generative AI, prompt engineering, agentic, MLOps, NLP, forward deployed) went from 0.20% in July 2024 to 0.50% in May 2026 (78,841 postings in May 2026) of US postings. The largest categories are machine-learning engineering, AI

eadership roles, AI and ML trainers, AI consulting, AI software development and AI product management.

5.9 The friction gap

Rated exposure (what a model could do) and observed exposure (what Claude is actually used for, from the Anthropic Economic Index, occupation level, March 2026 release) correlate 0.63 across 876 occupations (Spearman 0.72). The gap is systematic by occupation group: observed usage runs far above the rating in computer and math, legal, sales, office and finance work, and far below it in engineering, transportation, protective service, production and healthcare delivery. The groups with the largest shortfall are the ones where errors are physical, licensed or supervised, which is the friction hypothesis (H3) in the design.

5.10 Sorting by residual hard tasks (H4)

Splitting occupations by the share of their O*NET tasks rated not exposed (E0) in the Eloundou task file: β = +0.01 (s.e. 0.10) in occupations with few unexposed tasks left and +0.25 (s.e. 0.14) in occupations with many; the interaction is −0.10 (s.e. 0.10). No evidence that occupations with few hard tasks left are losing juniors faster.

Part II. The pre-ChatGPT window: July 2019 to May 2026

Part I had no baseline: postings began in June 2024, after ChatGPT, GPT-4 and the first coding assistants. Five further Revelio pulls (fiscal years July 2019 to June 2024, full schema) extend the panel to 2.53 billion postings and 83 months, 40 of them before ChatGPT. Not all of

. is usable for every series. Revelio's LinkedIn feed is under 1% of its US postings before May 2020 and ramps up, white-collar occupations first, between November 2020 and July 2021, so a LinkedIn-only contrast between exposed and unexposed occupations in those months measures LinkedIn's coverage, not hiring. LinkedIn-sourced series therefore start in July 2021 (16 months before ChatGPT), the first month from which LinkedIn's count stays above a quarter of its Jan to Jun 2022 level. Non-LinkedIn sources (Indeed and other aggregators before 2021; company sites and staffing firms join later) run from July 2019 and carry the COVID shock, which is used as a placebo: a large demand shock to the same occupations that had nothing to do with AI. The design changes accordingly: every outcome is now a monthly event study, occupation by month, with the average of July to October 2022 (the last four months before ChatGPT) as the reference, occupation and month fixed effects, Jan to Jun 2022 posting weights, and standard errors clustered by occupation. The coefficient in each month is the difference between a fully exposed and a fully unexposed occupation (Eloundou GPT-4 beta 0 to 1) relative to that difference in the reference months. Grey dotted lines are the major model releases; orange lines are the Fed hiking cycle and the big-tech layoff wave.

> **Key to the numbered markers (the numbers are the same on every chart in this paper):** (1) ChatGPT, Nov 2022 · (2) GPT-4, Mar 2023 · (3) o1-preview, Sep 2024 · (4) Claude 3.5 Sonnet, Oct 2024 · (5) DeepSeek-R1, Jan 2025 · (6) Claude 3.7 Sonnet, Feb 2025 · (7) Gemini 2.5 Pro, Mar 2025 · (8) o3 and o4-mini, Apr 2025 · (9) Claude 4, May 2025 · (10) GPT-5, Aug 2025 · (11) Claude Sonnet 4.5, Sep 2025 · (12) Gemini 3 Pro, Nov 2025 · (13) Claude Opus 4.5, Nov 2025 · (14) Claude Opus 4.6, Feb 2026 · (15) GPT-5.4, Mar 2026 · (16) Claude Fable 5 / Mythos 5, Jun 2026. Orange vertical lines are macro events (COVID emergency, Omicron wave, Fed hikes begin, big-tech layoff wave, Fed first cut). On charts of seasonal series, the grey bands step darker through the year (Q1 unshaded, Q4 darkest), so the shading resets at each new year.

II.1 Volume

II.2 Seniority

II.3 Pay

II.4 Robustness

Difference-in-differences coefficients (exposure x post-period, pre-period from July 2021), with the pre-trend (exposure x years, July 2021 to October 2022, same fixed effects) alongside. "All sources" pools LinkedIn and non-LinkedIn postings and is shown to document the coverage problem, not as evidence. The last column is the COVID placebo on the all-source panel (pre Jul 2019 to Feb 2020, post Mar 2020 to Feb 2021), which before July 2021 is the non-LinkedIn feed alone, LinkedIn months before then being masked.

> How to read Part II against Part I. Part I (June 2024 onward) found exposed occupations gaining postings relative to others. Part II shows why: they were recovering from a 2023 trough that Part I could not see. Over the full window the exposed-occupation gap opened in early 2023, held through 2025 and closed in 2026. The junior-share results are unchanged in substance: no measure shows exposed occupations shedding junior postings relative to senior ones in a way that survives within-group controls. The pay result is unchanged in direction and larger.

Part III. Headcount: Revelio workforce dynamics, Sep 2020 to Sep 2026

Parts I and II measured hiring intent. Part III measures people: Revelio's workforce-dynamics table gives, for every firm and month, headcount and inflow by seniority level (1 = entry to 7 = executive), built from professional profiles. Junior is seniority 1 to 2, senior 4 to 7. Firm AI exposure is the posting-weighted Eloundou rating of the firm's own US postings before ChatGPT (pre-ChatGPT postings (2019-07 to 2022-10), >= 30 matched US postings). The design is the same as Part II, at firm level: event studies relative to the Jul to Oct 2022 average with firm and month fixed effects, Jan to Jun 2022 headcount weights, standard errors clustered by firm; then a triple difference with firm x band, firm x month and band x month fixed effects; then the same within Revelio's role families so that a firm's mix of roles cannot drive the result. Descriptively, the junior share of headcount at the most exposed third of firms went from 67.9% in Sep 2020 to 69.4% in Sep 2026; at the least exposed third from 86.4% to 86.2%.

> Caveats specific to this table. Revelio headcount is inferred from profiles and scaled to firm size; seniority is Revelio's model, not a payroll grade; inflow counts positions started, which includes internal moves; firms without a measurable US posting mix before ChatGPT are excluded, which drops small and private firms. Coverage of the table itself (the number of firms and the scaled totals) is checked month by month in the log.

---

6. What this does not show

- **Power.** The pay-band β has a standard error of about 0.05 and the title-tier β about 0.28. The Canaries result translates to roughly −0.05 per year on this scale. The pay-band interval (−0.01 to +0.18) excludes it; the title-tier interval does not. The cleaner headcount test is the less powerful one.
- **The pre-ChatGPT window is not a clean baseline.** Part II has 16 months before ChatGPT in the LinkedIn-sourced series (40 in the other sources, which also carry the COVID shock), but exposed occupations were converging down from a 2021 hiring boom through the 2022 hiking cycle, and the big-tech layoff wave began three weeks before ChatGPT. Within-group estimates net out the tech cycle at the 2-digit level only. The pre-2024 pulls carry Revelio's seniority field; the 2024 to 2026 pulls do not, so the direct seniority test stops in June 2024 until those years are re-pulled with the full schema.
- **Postings are intent, not hires.** A firm can post a junior req and fill it with an agent, or post fewer reqs with larger expected hires. Expected-hires weighting is a planned robustness check.

Pay bands are a proxy for rung, and they move with pay. They validate where ladders are steep (software, law) and not where they are flat (analysts, clerical), and a fall in offered pay moves postings into the low band with no change in who is hired. Title tiers are immune to that but cover 15% of postings, 145 occupations with full coverage, and are seasonal.

- **Coverage drift.** Revelio's capture changed within the window. Source-interacted fixed effects remove level shifts by occupation and month; they cannot remove a source change that is itself tilted by tier within an occupation.
- **Exposure is static.** Eloundou ratings are a 2023 snapshot; the AEI observed measure is from 2026. Neither dates when a task crossed the frontier. The design in the original proposal (an agent battery re-run across model generations) would supply crossing dates; it is Part II.
- **Correlational extensions.** Levels, layoffs, wage offers and cross-country results share the sector-cycle confound and are labelled accordingly.

---

7. What comes next

The first version listed five items. Two are done and folded in above: the pre-ChatGPT window with the seniority field through May 2026 (Part II), and the firm-level employment test on the workforce-dynamics table (Part III). Still open:

1. Re-pull the nine source flags for July 2024 to June 2026 (a column pull joined on job_id, like the seniority supplement), to name which feed drove the 2024 to 2026 coverage expansion.

2. Vacancy-duration outcomes from the remove-date field, which now runs July 2019 to June 2026 (the 2024 to 2026 postings carry it on 95 to 99% of rows, via the job_id supplement), with expected-hires weighting; time-varying exposure from successive Anthropic Economic Index releases.
3. The outage design from the companion note (API outages as hourly withdrawal shocks on GitHub activity), which needs no permissions and is the cleanest causal idea in the agenda.
4. The agent battery, as a calibration appendix across 10 to 15 occupations rather than the spine of the paper.

Reproducibility

Each figure also exists as an interactive Plotly version, with a tidy CSV of its plotted data and a machine-readable spec, so every chart can be re-rendered in any framework; the bundle is published alongside this paper. Version history: the June 2024 window (Part I, first version, November 5, 2026); the July 2019 extension and the COVID placebo (November 8); the seniority field through May 2026 and the all-releases figure (November 9).

Pipeline: DuckDB streaming aggregation (stage 1: CSV to parquet; stage 2: seven aggregate panels), then Python for fixed-effects estimation (alternating-projection demeaning, cluster-robust sandwich). Public inputs: Eloundou et al. (2024) occupation and task files from the OpenAI GPTs-are-GPTs repository; Felten, Raj and Seamans AIOE; Anthropic Economic Index releases (Hugging Face); model release calendar compiled from announcements. Occupation-level aggregates and code will be released; firm-level series with company identifiers are held back pending the Revelio redistribution terms.