A Market for Human Expertise: Observations*

Alberto Cavallo Zoë Cullen

This draft: August 21, 2026

Abstract

The market for human expertise to train AI is among the fastest-growing marketplaces and cited as a primary source of new work created by AI, with potential to contribute 19B to GDP by 2030.[1] We use daily data from a leading platform where multiple AI labs purchase expertise, covering 134,000 hires and over $330 million in expected pay out at posted rates between November 2025 and August 2026, to ask how the market for expertise-as-training-data compares to the market for expertise-as-labor. We find evidence that new hires and wages for expertise-as-training-data in a domain fall as the accumulated corpus of expertise grows. Demand arrives in successive waves: most postings attain the majority of their eventual hiring in our window within four weeks of first posting and then new hiring subsides. On average, wages globally start above the US full-time equivalent for the same narrow occupation, and decline over time. The sequence of hiring waves is predictive of the closing of the AI adoption gap between observed and potential AI use in the occupation.

Data Description

We examine the complete public job board of a leading platform where AI labs buy expertise daily between November 2025 and August 5, 2026. Each snapshot records the job title, the pay, a persistent posting identifier, and a counter of the form "N hired this month." The counter is a rolling sum: it reports the posting's hires over the trailing 31 days. Day-over-day change incorporates both the current day's new hires with the exits from 31 days earlier. We recover daily hires by adjusting the counter's daily change for these exits.[2] Summing across the panel's

*Cavallo: acavallo@hbs.edu, Cullen: zcullen@hbs.edu

[1]The Economic Impact of the Data Annotation Industry, Oxford Economics, December 2025, url = https://www.oxfordeconomics.com/wp-content/uploads/2025/12/OxfordEconomics_DataAnnotationImpacts.pdf, accessed Aug 20, 2026.

[2]Appendix A states the recursion and its starting conditions.

1,962 postings gives roughly 134,124 hires with estimated compensation of $330 million at posted rates.[3]

The platform classifies 80% of jobs into 12 occupational domains, and we further map each job to the 2018 Standard Occupational Classification, using the detailed SOC-6 descriptions and linguistic similarities to the job descriptions.[4] When we aggregate to SOC 2-digit major groups, 88 percent of postings are assigned. We also collect the following details from each posting: education and experience requirements, the stated time commitment, and the form of compensation: a time rate (hourly, weekly, or yearly) or a fixed payment (one-time or per task), together with the range or amount specified and whether the posting's text offers a cash bonus or merit-based escalation of pay. We observe each posting's location terms: whether a location is required, the eligible or excluded countries, and whether the work is remote or on-site. Finally, we merge in occupation-level education and experience requirements from the Occupational Information Network (O*NET) and occupation-level wages and employment from the BLS Occupational Employment and Wage Statistics (OEWS).

For AI-use measures we use the occupation-level series of the Anthropic Economic Index, released March 5, 2026. *Observed AI use* is the share of each occupation's task time observably performed with AI, constructed by Massenkoff and McCrory (2026) from Anthropic usage logs. *Theoretical AI use* is the share of an occupation's tasks that frontier models could plausibly perform, taken from the human-rated full-exposure series of Eloundou et al. (2023). The *AI adoption gap* is theoretical minus observed use, i.e., the portion of feasible AI use that has not yet materialized.

Obs. 1: Hiring proceeds in successive waves for domains of expertise.

Figure 1 plots daily new hires for the work domains on the platform. Each series is scaled to its own peak, with domains ordered by the hires-weighted mean date of hiring, putting the latest wave at the top. The platform labels twelve work domains, with eleven exceeding 1,000 hires over the nine-month panel. Two features of these data are worth emphasizing.

First, within most domains, hiring arrives in waves, rising steeply, peaking, then declining. Across postings, the majority of eventual hiring over the nine-month window occurs within the

[3] The compensation estimate combines posted wages with reported or estimated job durations. Survey postings, which account for 4% of hires, specify 15- to 60-minute engagements and pay per response, which we use directly, as we do for other task-based and one-time-priced postings. For hourly postings, we use stated weekly hours where available (one-third of postings) and assume 40 hours otherwise, over a one-week engagement. Appendix C.1 details the calculation and Table 1 reports the assumptions. Assuming 40 hours for every hourly hire gives $363 million.

[4] 67 percent of hires are on titles coded to a detailed (6-digit) occupation, 4 percent only to a coarser 3-digit minor or 2-digit major group, and 29 percent (generalist and annotation work) have no occupational home and are excluded from occupational aggregates.

first four weeks of the posting's first appearance. After that, hiring falls to 9% of the peak weekly rate, on average, and 33% among postings still on the board.

This pattern is consistent with the notion that once a corpus of expertise is assembled, the marginal value of an additional expert falls, at least in the near term. One important caveat, however, bears on this interpretation: the working relationship can continue indefinitely, producing new training data without showing up as a new hire in our panel.

Second, waves across domains arrive successively. Language and transcription comes first and is largely finished before the others begin. Software engineering and data analysis follow. Demand then shifts toward professional services. Finance, law, medicine, and business operations all reach 50 percent of their hiring within a five-day span. The sciences sit between these groups and show the least wave-like profile: life, physical, and social science peaks at 255 hires per day on April 21 yet is still running at 70.7 percent of that peak several months later at the panel's end. In Obs. 4, we show evidence that the early hiring waves are associated with greater AI adoption relative to the technology's potential.

These wave patterns are robust to alternative constructions and classifications. Appendix Figure 8 shows that the waves are present in raw daily hires without the 7-day smoothing, while Figure 9 shows smoother but similar waves when we convert hires into the stock of experts at work using posting-level duration assumptions. Figure 7 shows that the pattern also persists under the SOC classification, even though these broader occupation groups can combine distinct types of expertise with different timing, such as Office and Administrative Support combining parts of data analysis and business operations.

Obs. 2: Wages start, on average, above the US full-time equivalent for the same narrow occupation.

Expertise-as-training-data is exchanged in a global marketplace, but the same experts can also supply their labor in traditional labor markets. We therefore compare posted wages on the platform with wages for the same occupations in the US labor market.

To construct this comparison, we map each job posting to a detailed (6-digit) SOC occupation using its title and description and match it to wages from the BLS Occupational Employment and Wage Statistics (OEWS). The latest available OEWS data refer to May 2025, so we express both sides in dollars of the same quarter, 2026Q2, using the Employment Cost Index for wages and salaries of private-industry workers.

Most listings for expertise-as-training-data post an hourly pay rate, with less than 10% offering contingent pay for completed tasks.[5] Among hourly postings, the hires-weighted average posted range is 27% of the midpoint, and 35% of hires face a single posted wage. We use the

[5] Contingent pay becomes increasingly common over the panel, rising to nearly half of new hires in the latest period.

midpoint of the posted range in Figure 2, but relative wages remain higher on average using the wage floor instead.

The hires-weighted average wage ratio is 1.47 times the US mean wage for the same narrow occupation. The premium appears in every domain: it is largest in Finance (2.09) and Business Operations (2.01), and smallest in Language and Transcription (1.04), which is also the largest and earliest hiring wave in Figure 1, with 21,221 hires. Posted rates also closely track the US wage structure across domains: domains that pay more in the US also pay more on the platform. The fitted relationship lies above the 45-degree equal-pay line but has a slope of 0.81, indicating that the premium narrows as the US wage of the occupation rises.

The result is robust to using our alternative SOC classification. Appendix Figure 12 groups listings by SOC-2 major group rather than by the platform's domains.[6] Although the ordering of premiums changes across groups, the relationship between platform and US wages remains positive and less than proportional ($\rho = +0.66$).

The premium remains when using the floor rather than the midpoint of each posted wage range. The mean ratio falls from 1.47 to 1.24, and the ranking of major groups remains similar, although the premium is no longer universal. Office and Administrative Support remains well above the US wage (2.25), followed by Business and Financial Operations (1.82), while Computer and Mathematical (0.92), Arts, Design, Entertainment, Sports, and Media (0.90), and Life, Physical, and Social Science (0.88) fall slightly below parity. The across-occupation relationship is also little changed ($\rho = +0.60$). Thus, even assigning each posting the bottom of its wage range, platform wages average about one quarter above the US mean for the same occupation.

The premium is not explained by the geographic composition of hiring. Although our benchmark is the US wage, only 8 percent of hires are on listings restricted to US workers, and US restrictions are not concentrated in the groups with the largest premiums (Figure 13). The three groups with the highest pay ratios—Office and Administrative Support, Sales, and Business and Financial Operations—restrict 15, 5, and 5 percent of hires to the US, respectively. By contrast, the two groups with the highest US-only shares, Management (32 percent) and Healthcare Practitioners (21 percent), have pay ratios of 1.45 and 1.15.

These comparisons nevertheless have important limitations. They concern posted cash compensation and therefore do not account for benefits associated with traditional employment, which could offset some or all of the observed wage difference. In addition, the OEWS mean pools all workers within an occupation, whereas the platform screens applicants for particular tasks, so worker composition may differ even within a 6-digit SOC occupation. However, Appendix C.2 shows that stated degree and experience requirements are lower on the platform than in the corresponding occupations, providing no evidence that the wage premium is driven

[6] The average pay ratio coincides between the domain and the SOC-2 figures because we weight listings rather than the categories after aggregating: each listing's ratio counts by the listing's hires, and the average is taken over listings before they are grouped, so regrouping them by SOC-2 major group leaves the 1.47 of the main text unchanged.

by more stringent observable credential requirements.

Obs. 3: Wages for expertise in a given domain decline over time.

We next examine how the price of expertise evolves over time. A simple average of posted wages moves with the composition of listings, not only with the rates attached to them. For example, if the platform shifts from recruiting physicians to recruiting medical transcriptionists, the average posted wage within the same domain might fall even when no individual listing has changed its rate. We therefore construct a matched-listing price index that follows the same listings over time. Figure 3 plots this index separately for each expertise domain, together with a platform-wide weighted average. To construct the aggregate, we hold the composition fixed, weighting each domain by its wage bill over the panel — its total hires $\times$ mean posted rate. The aggregate index therefore isolates changes in posted rates from shifts in the composition of expertise across domains.

We find that the price of expertise falls over our sample. In the aggregate, the matched-listing index declines by 10.1 percent. The largest declines occur in Medicine and Language and Transcription, where matched-listing wages fall by roughly 22 and 19 percent, respectively, while rates in several other domains remain comparatively stable. This decline is especially notable because posted wages for continuing listings are relatively sticky, with adjustments occurring as infrequent discrete changes.

Wage adjustment may also occur through the rates attached to newly created listings. In Figure 4, we construct an index based on the initial hourly-equivalent wages posted by newly appearing listings. This index is considerably more volatile than the matched-listing index, suggesting that new listings may provide an additional margin through which compensation adjusts. However, comparisons across new listings do not hold the underlying expertise fixed. Listings entering at different points in time may differ in specialty, qualifications, seniority, task requirements, or other characteristics that we do not observe. We therefore treat the new-listing evidence as suggestive of an additional adjustment margin rather than as a price index for fixed-quality expertise.

Obs. 4: The timing of hiring predicts the AI adoption gap and AI use.

The timing of hiring for expertise in a specific occupational domain predicts the AI adoption gap in that occupation, constructed by Massenkoff and McCrory (2026) and building on Eloundou et al. (2023). Figure 5 plots the timing of each group's hiring wave (the hires-weighted mean date) against its AI adoption gap, as measured in March 2026. Wave timing, weighting each major group by its hires, and the March 2026 AI adoption gap are positively correlated ($\rho = +0.41$).

In contrast, AI use is negatively correlated with wave timing ($\rho = -0.30$) and potential use for automation exhibits zero correlation, as shown in Figure 6. The most recent waves may have targeted areas where AI's usefulness lagged what the technology could feasibly support, and the early waves may have contributed to the higher rates of AI use and applicability. The total size of the marketplace, measured both in hires and estimated total spend, correlates most strongly with AI's overall potential for automation, as reported in Appendix Figure 18. The result is also robust to restricting to the eleven SOC-2 groups with at least 1,000 hires over the panel, the selection used elsewhere in the paper, where the hires-weighted correlation is effectively unchanged (Figure 19).

Conclusion

The market for human expertise used to train AI can differ from a traditional labor market in an important way: expertise purchased today may reduce the value of purchasing more of the same expertise tomorrow. We find patterns consistent with this hypothesis. Experts are initially hired at high posted wages, but hiring in many domains arrives in concentrated waves and then declines, while wages for continuing listings also fall. Demand then shifts across domains, with later waves concentrated in occupations with larger gaps between potential and observed AI use.

Together, these patterns suggest a market in which human expertise is especially valuable while AI capabilities are being built, but where the accumulation of training data can reduce the marginal value of additional expertise in the same domain. This process is not universal: in areas such as the sciences, hiring and wages remain comparatively persistent. Our evidence covers nine months and a single large platform, but it points to a distinctive lifecycle of AI-related expert work: high initial demand and compensation followed, in some domains, by declining demand as the stock of relevant training data grows.

References

[1] Eloundou, Tyna, Sam Manning, Pamela Mishkin, and Daniel Rock. 2023. "GPTs are GPTs: An Early Look at the Labor Market Impact Potential of Large Language Models." Working paper.

[2] Massenkoff, Maxim, and Peter McCrory. 2026. "Labor Market Exposure to AI: Theory and Evidence from Anthropic Usage Data." Working paper. Data released March 5, 2026, Anthropic Economic Index.

[3] Cullen, Zoë B., and Nicole Tempest Keller. 2026. "Humans in the Loop (HITL) AI." Harvard Business School Background Note 826-218.

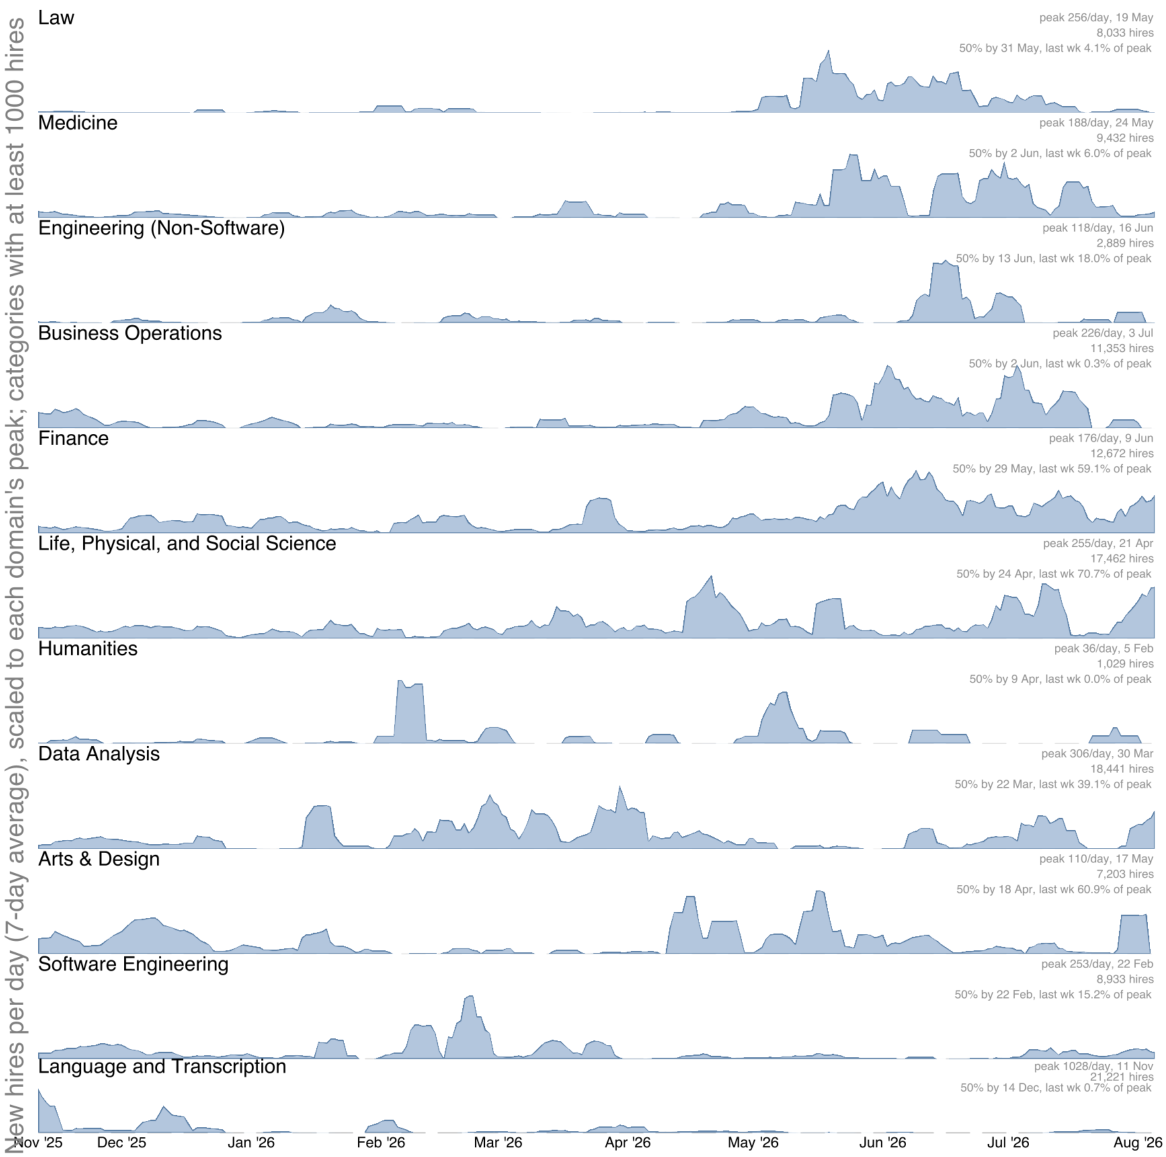

Figure 1: Hiring waves by the talent-sourcing platform's own listing-domain tag, categories with at least 1,000 hires. Rows ordered by the hires-weighted mean hiring date. Right-hand labels report each domain's peak daily hiring rate and its date, total hires, the day its cumulative hiring reaches half of its window total, and the final week's average daily rate as a share of the peak.

Figure 2: Relative Wages by Expertise Domain

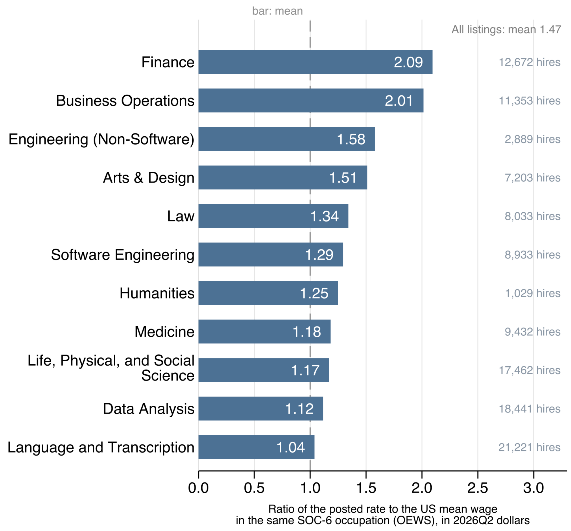

A. Mean pay ratio by domain: posted midpoint over the mean OEWS wage in the same narrow occupation (SOC-6), weighted by hires

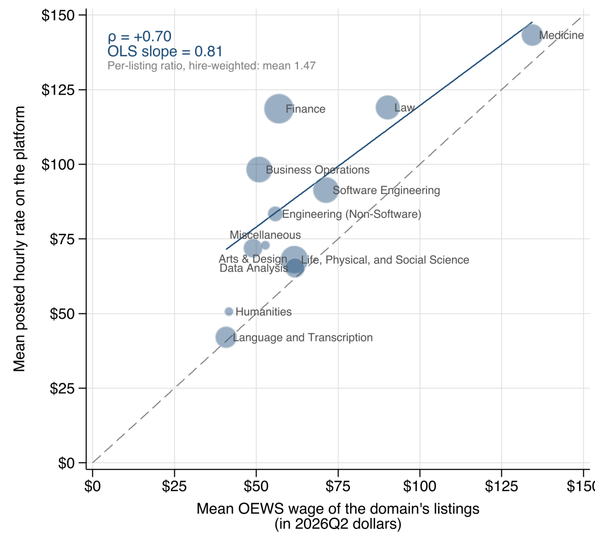

B. Mean posted midpoint against the mean OEWS wage of the domain's listings (matched at SOC-6), weighted by hires

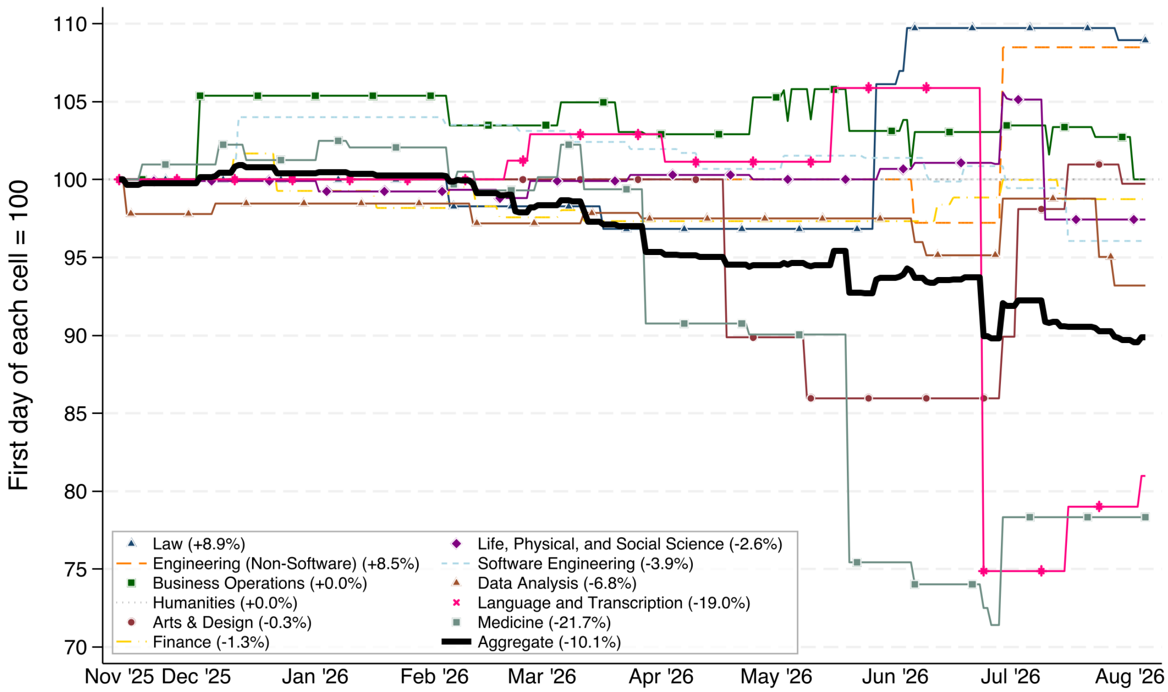

Figure 3: **Matched-listing wage indices by expertise domain.** The figure reports matched-listing wage indices separately by expertise domain, together with the platform-wide weighted average. Each index follows continuing listings over time, so changes reflect revisions to posted rates rather than the entry or exit of listings. The plotted domains are the eleven with at least 1,000 hires, the sample of Figure 1. The aggregate holds the talent-sourcing platform's composition fixed, weighting each domain by its wage bill over the panel (total hires times the mean posted rate), and runs over all twelve domains the platform labels, including the Miscellaneous catch-all that the panels leave out.

Figure 4: Wage rates on continuing and newly listed jobs

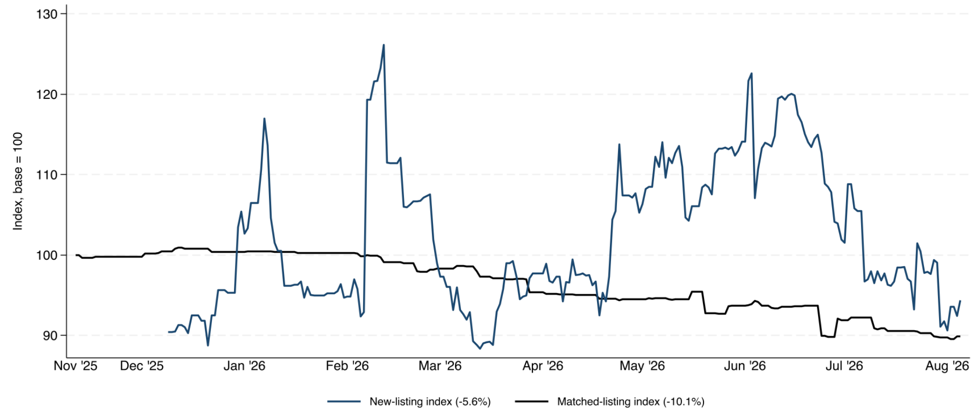

Notes: The matched-listing index compares posted hourly-equivalent rates for continuing listings and aggregates changes across fixed expertise-domain cells. The new-listing index measures the initial posted rate of listings first appearing within the trailing 28 days, aggregated using the same fixed cells. Both aggregates hold the domain composition fixed, weighting each domain by its wage bill over the panel (total hires times the mean posted rate).

Figure 5: Timing of expert hiring and the AI adoption gap

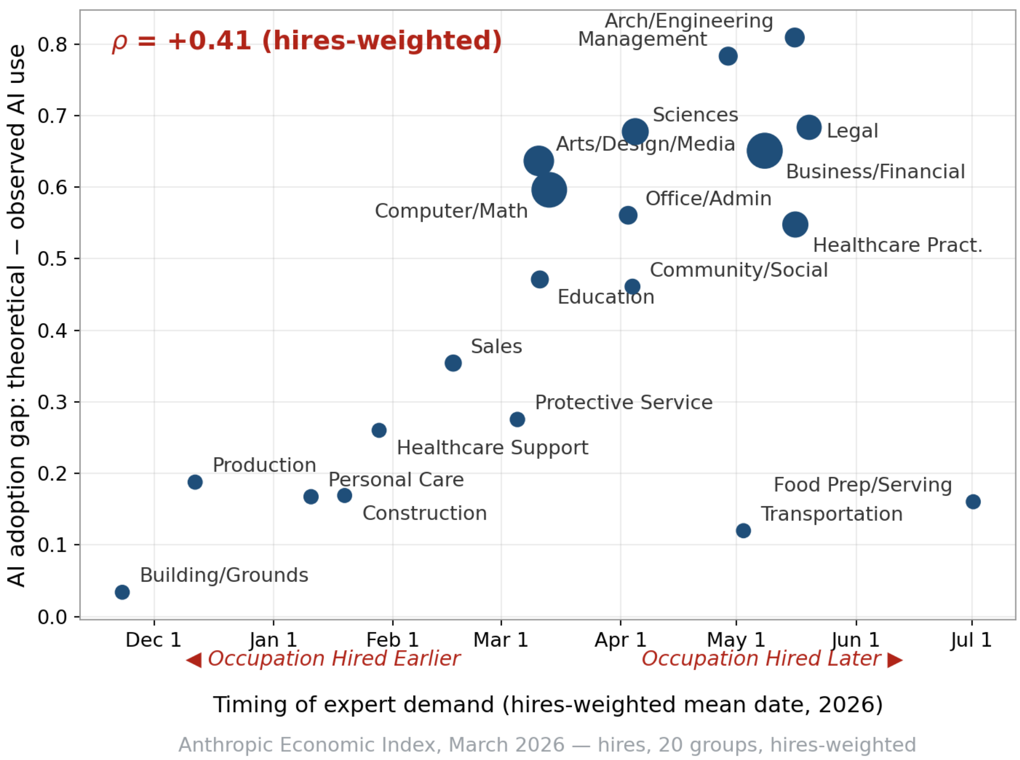

Notes: Each point is one of the twenty SOC 2-digit major groups that appear on the board (Data Description). The horizontal axis is the timing of the occupation group's hiring through the platform: the hires-weighted mean date over November 11, 2025 to August 5, 2026, with daily hires recovered as in Appendix A. The vertical axis is the group's AI adoption gap: theoretical minus observed AI use, employment-weighted (Anthropic Economic Index, March 2026). Marker area scales with the group's total hires. The ρ weights each group by that same quantity, so a group's influence on the statistic is what its marker shows. The mean dates of the groups still hiring at the August 5 end of the panel (arts/media, business/financial, and healthcare practitioners, each with at least 10 percent of their hires in the final 28 days) are right-censored, so their true dates are later than shown, which likely biases the plotted correlation downward.

Figure 6: Correlation of hiring-wave timing with AI-use measures

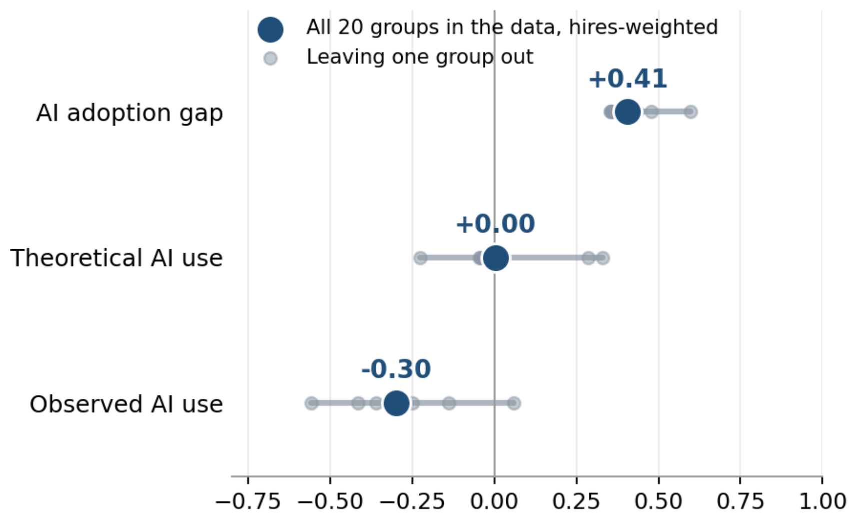

Notes: Correlation between the timing of each group's hiring wave (the hires-weighted mean date of the platform's hiring, November 11, 2025 to August 5, 2026, with daily hires recovered from the trailing counter as in Appendix A) and three occupation-group AI-use measures, across the twenty SOC 2-digit major groups that appear on the board, with each correlation weighted by the group's hires. All three measures are employment-weighted aggregates over each group's occupations from the Anthropic Economic Index release of March 5, 2026: observed AI use from Anthropic usage logs, theoretical AI use from the human-rated full-exposure series of Eloundou et al. (2023), and the AI adoption gap, their difference. Large dots use all twenty groups; small dots re-compute each correlation leaving out one group at a time, and the bar spans the resulting min–max range. Timing correlates +0.41 with the adoption gap and stays positive in every leave-one-out sample (+0.35 to +0.60). It is uncorrelated with theoretical use on its own and negatively correlated with observed use (−0.30): what dates a domain's wave is not how automatable the work is, nor how much AI is already used there, but the distance between the two. The mean dates of the groups still hiring at the August 5 end of the panel (arts/media, business/financial, and healthcare practitioners) are right-censored, which likely biases the plotted correlations toward zero.

Appendix

A Daily hires from the trailing counter

Each posting displays a counter of the form "N hired this month." The counter is a rolling 31-day sum: on day t it shows $c_t = \sum_{s=t-30}^{t} h_s$, where h_s is the number of hires the posting made on day s. Its daily change is therefore

$$c_t - c_{t-1} \;=\; h_t - h_{t-31},$$

the current day's hires minus the hires that fall out of the window. We recover the daily series by inverting this identity, $h_t = (c_t - c_{t-1}) + h_{t-31}$, and bound each h_t between 0 and c_t so that measurement noise cannot produce negative hires or more hires than the counter displays. When a posting skips a day on the board, we carry its last counter forward, so a multi-day change is attributed to the day the posting reappears.

The recursion requires starting values, which depend on when a posting first appears. For postings that enter during our panel, we treat the first observed date as the posting's start date. The counter shown on that first day is therefore assigned to hires on that day. For the next 30 days, no hires can yet be falling out of the 31-day window, so daily hires are identified directly by changes in the counter. Beginning on day 31, we apply the recursion above, adding back the hires from 31 days earlier.

Postings that are captured on the first day of the panel require a separate initialization. There are 277 such postings, representing 14 percent of postings, with day-one counters totaling 7,361 hires, or 5.5 percent of all hires. Because these counters summarize hiring over the preceding 31 days, we cannot date those hires individually. We assume that each posting's initial counter accrued evenly over that 31-day window. Thus, one thirty-first of the initial counter is assigned to the first panel day, while the remaining shares are treated as pre-panel hires. These imputed pre-panel values provide the initial lagged hires needed to start the recursion. Because subsequently recovered hires themselves enter the add-back term 31 days later, the initialization can affect the estimated daily series beyond the first month and, in principle, throughout the panel.

B Hiring Waves Across Classifications

Regardless of how we classify hiring and work, our evidence suggests that the nature of expertise-as-training-data appears as waves.

Figure 7 redraws Figure 1 with hiring classified by SOC 2-digit major group (the title-to-occupation mapping of the Data Description) instead of the platform's own domain tags. The waves survive the reclassification, and corresponding categories peak on the same days: Legal peaks on May 19 as Law does, Healthcare Practitioners on May 24 as Medicine, Architecture and Engineering on June 16 as Engineering (Non-Software), and Computer and Mathematical on February 22 as Software Engineering.

Figure 8 instead removes the 7-day moving average from the domain figure. The rises, peaks, and declines are in the raw daily series. Figure 9 recasts the flow of new hires as the stock of experts at work under the lengths the postings state; Section C.1 explains how those durations are calculated, and assumptions made when the listing does not include any details about the length. Across these alterative visualizations of the hiring and work carried out, successive waves across domains appear.

C Expertise-As-Training-Data Contracts

C.1 Work Duration

The board's structured fields do not include the length of engagement. Hence, we have a large language model read the full description of each of the 1,962 postings and report what the posting says about the engagement: how long the work lasts and how many hours a week it takes.

Panel A of Table 1 reports the results of this process. Most postings do not include an end date. The 352 postings that fix an end carry 19.3 percent of hires, and the median length they state, weighted by hires, is 24 days. For Figure 9, we assert that postings without a specified end date last for the median duration of those that so assert an end date.

Panel B repeats the exercise for work intensity: how many hours a week the work requires. Half of the postings include details about the number of hours per week, 989 listings carrying 56.2 percent of hires. The median stated intensity is 20 hours a week, weighted by hires. When calculating the total spend, we include intensity and the assumptions listed in Panel B of Table 1.

Figure 9 turns the stated lengths into a stock: each dated hire occupies the days of its posting's engagement, and the curve counts the experts at work each day (the weekly hours matter only when hourly rates are converted into pay). The hiring waves of Figure 1 reappear as waves of experts at work.

C.2 Degree and Experience Requirements

We map job descriptions to their detailed O*NET-SOC classifications. This allows us to compare the degree and experience requirements observed among workers in these narrow SOC 6-digit occupations. Overall, we find that the required degrees and years of experience are lower in the listings for expertise-as-training-data. Figures 10 and 11 report the gaps by SOC-2 major group: each listing is compared with the median incumbent of its own SOC-6 occupation in the O*NET education and training questionnaire[7]: required schooling in years in one figure, required years

[7] We benchmark against the median rather than the mean answer because each SOC-6 occupation's distribution comes from a limited number of surveyed incumbents: with few respondents the mean is noisy, while the median

of related work experience in the other, and the per-listing gaps are averaged within each group with hires as weights.

On average a listing asks for 1.9 fewer years of schooling and 1.8 fewer years of experience than the median incumbent of its occupation. What's driving this, however, is the absence of degree and experience requirements in the marketplace for human expertise-as-training-data. Among the listings that do state a requirement (panel B of each figure), the gaps nearly close, to -0.1 years of schooling and -0.2 years of experience. One potential explanation for this is that the screening for expertise through AI interviews on the platform allows the platform to screen for expertise best matched for the specific task.

C.3 Wage Rates

Using the mapping between job postings and the detailed occupation (SOC-6), we can compare the wage and requirements offered in this marketplace to the wages and requirements in more standard labor markets, using the OEWS from May 2025. The wage we take from each posting is the last rate it shows, so the two sides are observed at different dates; we express both in dollars of the same quarter, 2026Q2, moving each from the quarter in which it is observed with the Employment Cost Index for wages and salaries of private-industry workers.[8]

On average, the posted midpoint is about 1.5 (1.47) times the US mean wage for the same occupation, despite being a global marketplace for talent. Figure 12 reports the construction behind that number. Figure 12A shows the ratio by SOC-2 major group: each listing's midpoint over the OEWS mean wage of its own SOC-6 occupation, averaged within the group with hires as weights, for the eleven major groups with at least 1,000 hires over the panel. Figure 12B shows the same comparison in levels, one point per SOC-6 occupation: the occupation's hire-weighted mean posted midpoint against its OEWS mean wage. For robustness, Figure 14 replaces means with medians throughout; in other words, each listing's midpoint over the OEWS *median* wage of its occupation, summarized by hire-weighted medians. Figure 15 repeats the exercise at the floor of the posted range; Figure 16 drops the hire weights, giving every listing an equal vote. Figure 2 of the main text keeps the SOC-6 match but aggregates by the platform's domains rather than by SOC-2: each domain's mean posted rate is plotted against the mean OEWS wage of its listings, both sides averaged over the same listings and weighted by hires within each domain. Figure 13 shows the hire-weighted share of listings restricted to the US in each SOC-2 group, with the rows in the order of Figure 12A: the groups with the highest pay ratios show low US-only shares, so the premium is not explained by listings that require US-based workers.

is more stable.

[8] Weekly, monthly, and yearly posted rates are converted to hourly equivalents at 40 hours per week and 2,080 hours per year (2,080/12 per month); OEWS occupations reporting only annual pay are converted at the same 2,080 hours. We use the not-seasonally-adjusted index; 2026Q2 is the latest quarter it covers.

C.4 Contingent Pay

Pay has shifted from hourly rates toward piece rates. The share of expert pay delivered as hourly compensation fell from 97.5% to 58.0% over the sample (Figure 17), with the balance moving to task-based rates and performance bonuses. Piece rates require verifiable output, and their spread indicates that labs have built the measurement infrastructure to identify the quality of expertise directly. Contingent pay and credentials appear as substitutes at the frontier. The piece-rate share is 7% higher in listings without explicit degree or experience requirements than in those stating at least one, and more than twice as high as in those stating both.

D Market size and AI Adoption Measures

We show how total hires, prices, and total estimated spend correlate with Anthropic's measure of observed AI use, theoretical potential of AI, and the gap between theoretical and observed AI use. Figure 18 reports those correlations; Figure 6 reports the same exercise for the timing of each group's hiring wave.

Figure 7: Hiring waves by SOC major occupation group

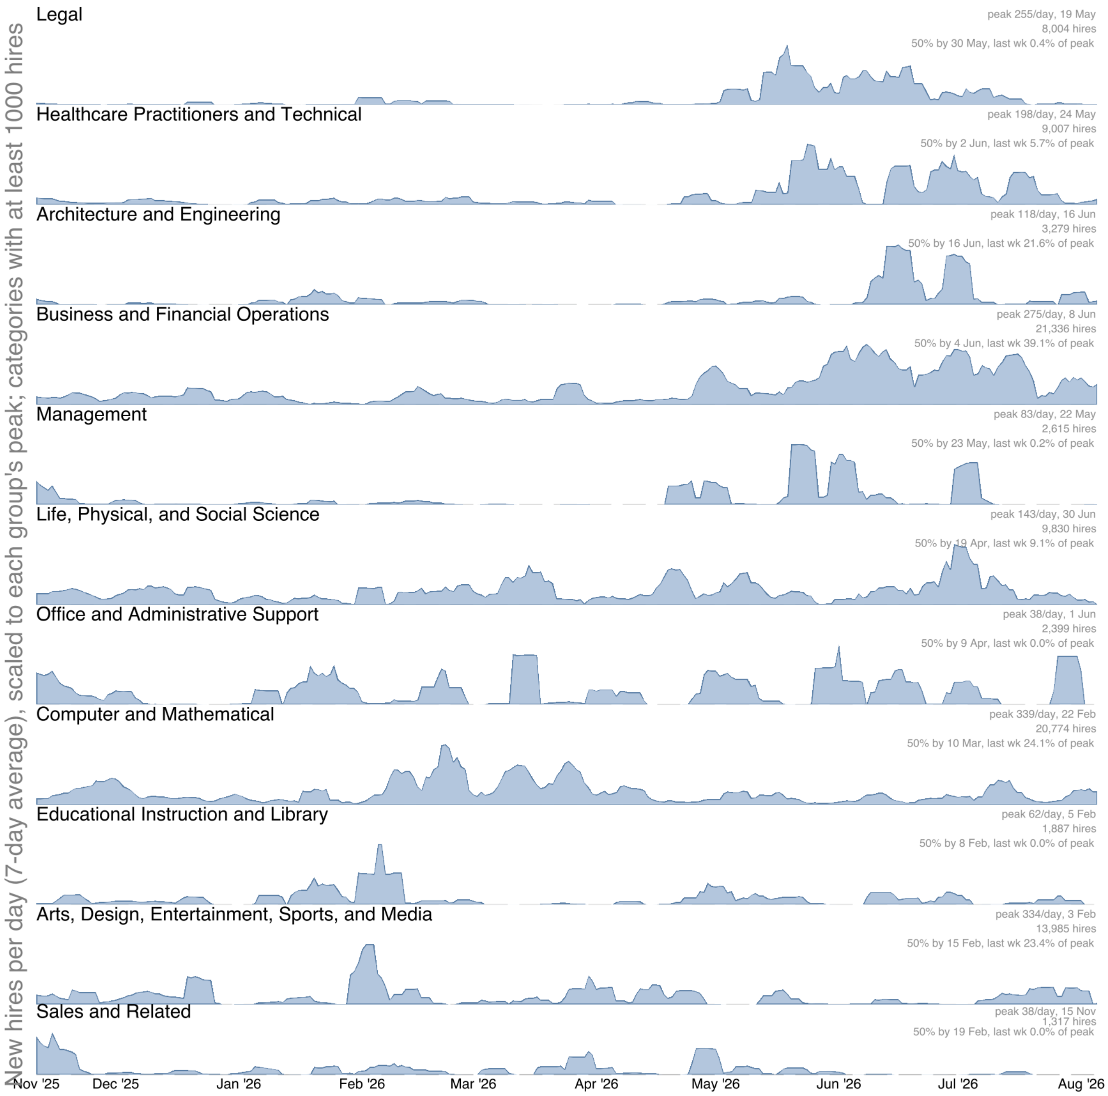

Notes: The platform's public job board, November 11, 2025 to August 5, 2026. Each row is a SOC 2-digit major group with at least 1,000 hires over the panel (eleven groups). Each curve is the group's new hires per day (7-day centred moving average), measured as posting-level day-over-day increments of the "N hired this month" counter, adding back the hires that disappear from the counter after a month, and scaled to the group's own peak; postings are assigned to detailed SOC occupations from their titles and aggregated to major groups. Hires recorded on a posting's first appearance count on that day; for postings already live on the panel's first date, the day-one counter is spread evenly over its trailing window (Appendix A), and the row totals include those undated day-one hires. The right-hand labels report the peak daily hires of the plotted series and the day it is reached, the group's total hires over the panel, the day the group's cumulative dated hires reach half of their window total, and the mean daily rate over the panel's final seven days as a share of the peak. Groups are ordered by the hires-weighted mean date of hiring, latest at the top. The top rows are right-censored by the August 5 end of the panel.

Figure 8: Hiring waves by expertise domain, without the 7-day moving average

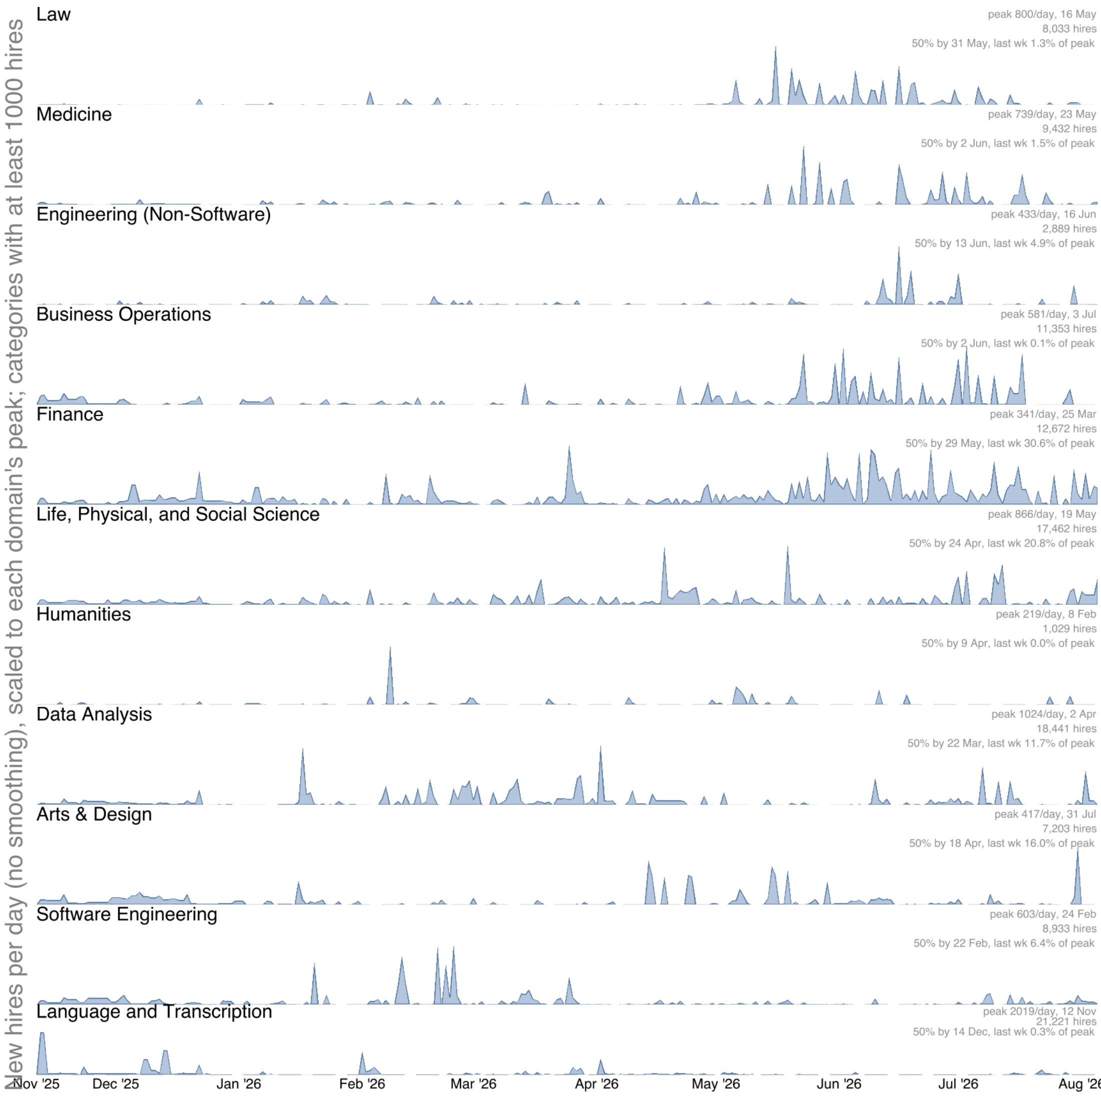

Notes: As Figure 1, without the 7-day smoothing: each row is the domain's raw daily hires, scaled to its own raw peak. The labelled peak level, its date, and the final-week share refer to this raw daily series; the day the domain reaches half of its hires does not depend on the smoothing.

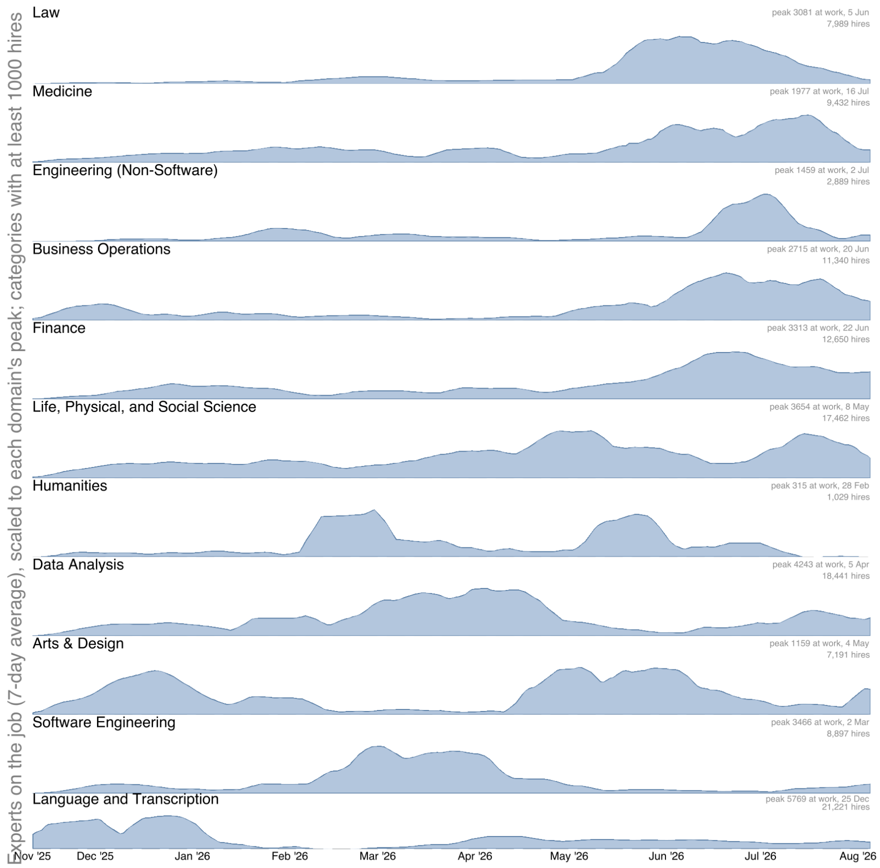

Figure 9: Experts on the job, at the lengths the postings state. Each dated hire occupies the days of its posting's engagement (Table 1, Panel A); the postings that fix no end take 24 days, the median among those that do, except those stating a floor, which keep the floor. Curves are 7-day averages of the experts at work each day, scaled to each domain's peak; rows, sample and order as in Figure 1.

Table 1: Work Duration and Intensity

What the posting says	Postings	Hires (%)	What we take
Panel A. How long the work lasts			
No end to the work			
Nothing about length	1,262	57.0	the median below, 24 days
Ongoing or long-term, no number	248	14.0	the median below, 24 days
A minimum ("at least four weeks")	100	9.7	the floor
An end to the work			
A single length ("three months")	164	6.7	the number
A range ("one to two months")	147	7.1	the midpoint
One unit of work ("about 30 minutes")	30	5.4	one day
Total hours only ("five hours total")	11	0.0	eight-hour days
States an end	**352**	**19.3**	**median 24 days (hire-weighted)**
Panel B. Hours per week			
No intensity stated			
Nothing about intensity	973	43.8	the platform's field, then a 40-hour week
An intensity stated			
A minimum ("at least 10 hours")	408	22.8	the floor
A typical load ("10 to 20 hours")	471	21.8	the midpoint
A ceiling ("up to 40 hours")	101	11.6	the ceiling
A daily amount ("8 hours per day")	2	0.0	five such days
An office schedule ("9:00 to 5:30")	7	0.0	a 40-hour week
States an intensity	**989**	**56.2**	**median 20 hours (hire-weighted)**

Notes: The cleaned job board, one reading per posting; a posting relaunched under a new listing is read from its first. Within each panel the rows partition the 1,962 postings, and the percentages are shares of hires, not of postings: the 1,262 postings silent about length are 64 percent of postings but carry 57 percent of hires. The bold row totals the block above it.

Figure 10: Required education against the occupation's incumbents, by SOC-2

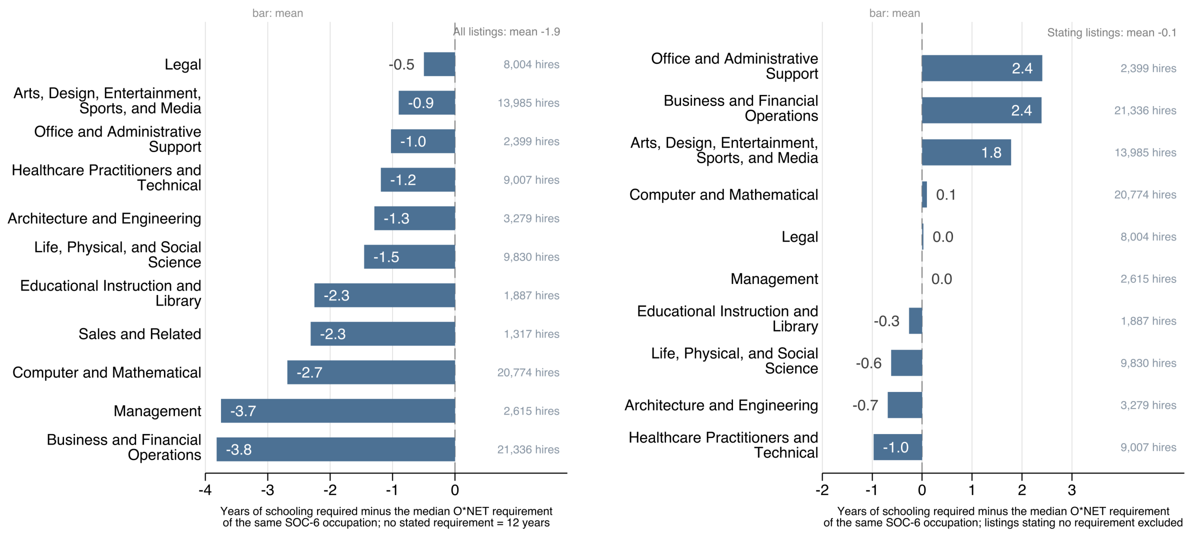

A. All listings, silence at the high-school base

B. Listings stating a requirement

Notes: Each bar is the group's hire-weighted mean gap, in years of schooling, between what the listing requires and the median incumbent of its SOC-6 occupation in the O*NET education and training questionnaire. In panel A a listing stating no requirement enters at the high-school base; panel B keeps only the listings that state one, against the same benchmark — Sales drops for lack of stating listings. Groups are the eleven SOC-2 majors with at least 1,000 hires; the corner statistic pools all listings, and the right column reports the group's board hires in both panels.

Figure 11: Required experience against the occupation's incumbents, by SOC-2

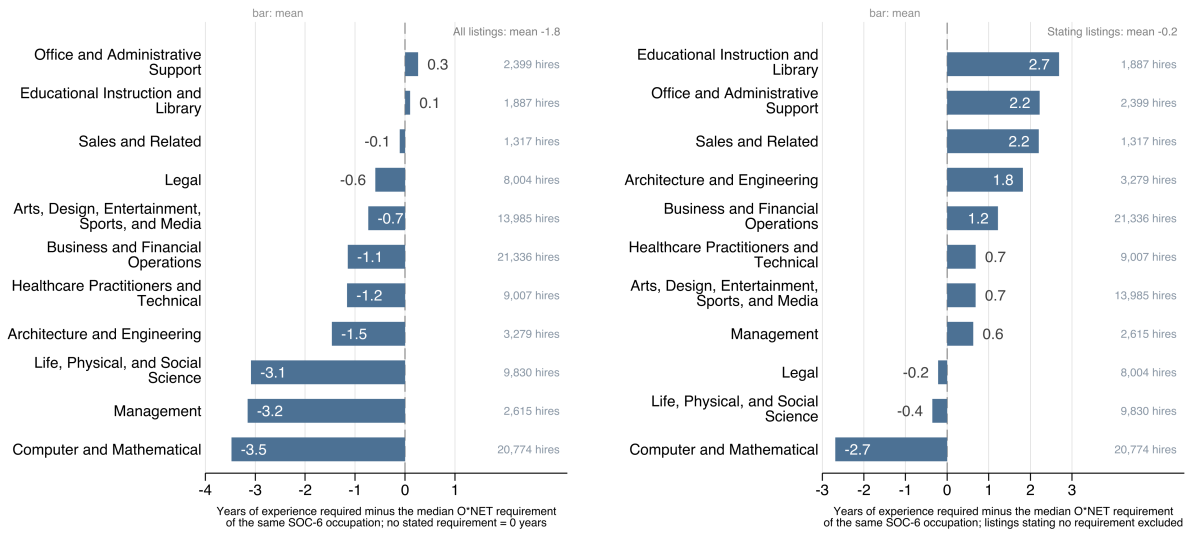

A. All listings, silence at zero years

B. Listings stating a requirement

Notes: Each bar is the group's hire-weighted mean gap, in years, between the experience the listing requires and the median incumbent of its SOC-6 occupation in the O*NET related work experience question. In panel A a listing stating no requirement enters at zero years; panel B keeps only the listings that state one, against the same benchmark. Groups are the eleven SOC-2 majors with at least 1,000 hires; the corner statistic pools all listings, and the right column reports the group's board hires in both panels.

Figure 12: Posted pay against the US wage, by occupation

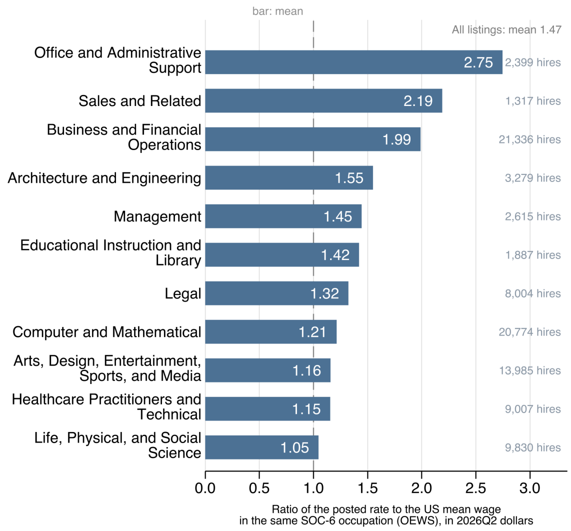

A. Mean pay ratio by SOC-2 (midpoint over OEWS mean), weighted by hires

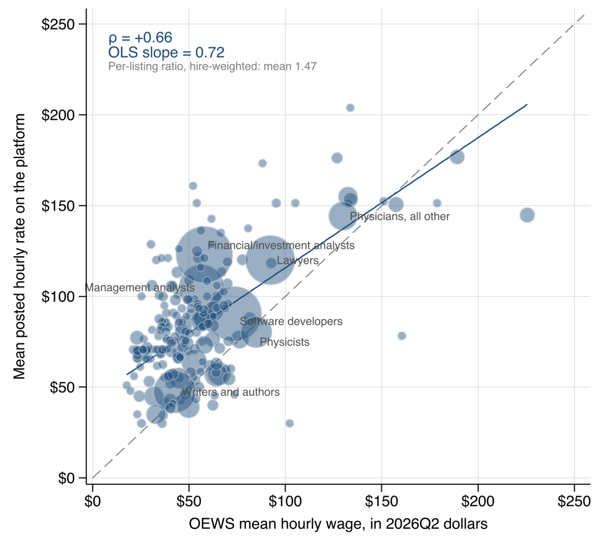

B. Mean posted midpoint against the US mean wage, by SOC-6, weighted by hires

Notes: In panel B each marker's area scales with the occupation's hires; the dashed line is equal pay, the solid line the hire-weighted least-squares fit. The corner statistic in both panels is the same hire-weighted mean ratio, computed over all listings before the 1,000-hire cut.

Figure 13: Share of listings restricted to the US, by SOC-2

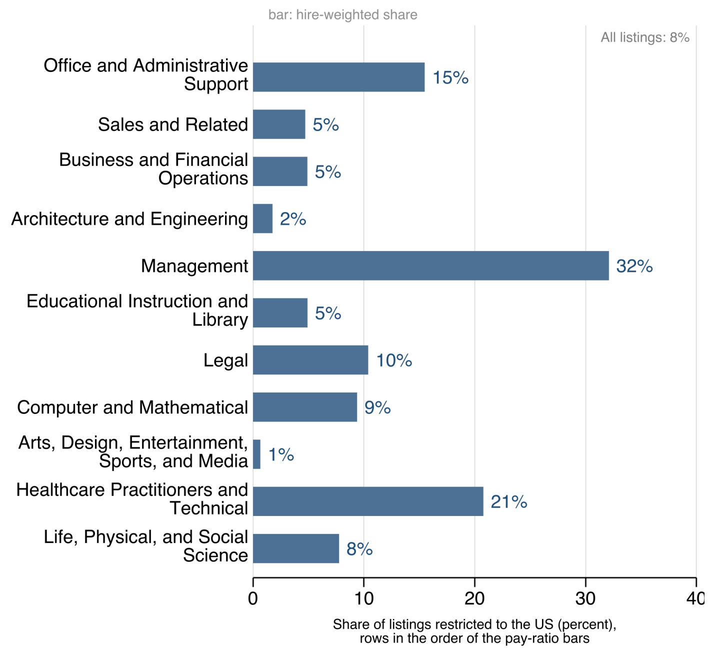

Notes: Hire-weighted share of listings whose required-country set is exactly the US, over the same listings as Figure 12A and with the rows in the same order — sorted by the pay ratio, not by the share — so the two figures read side by side. The corner statistic pools all listings before the 1,000-hire cut.

Figure 14: Robustness: medians in place of means

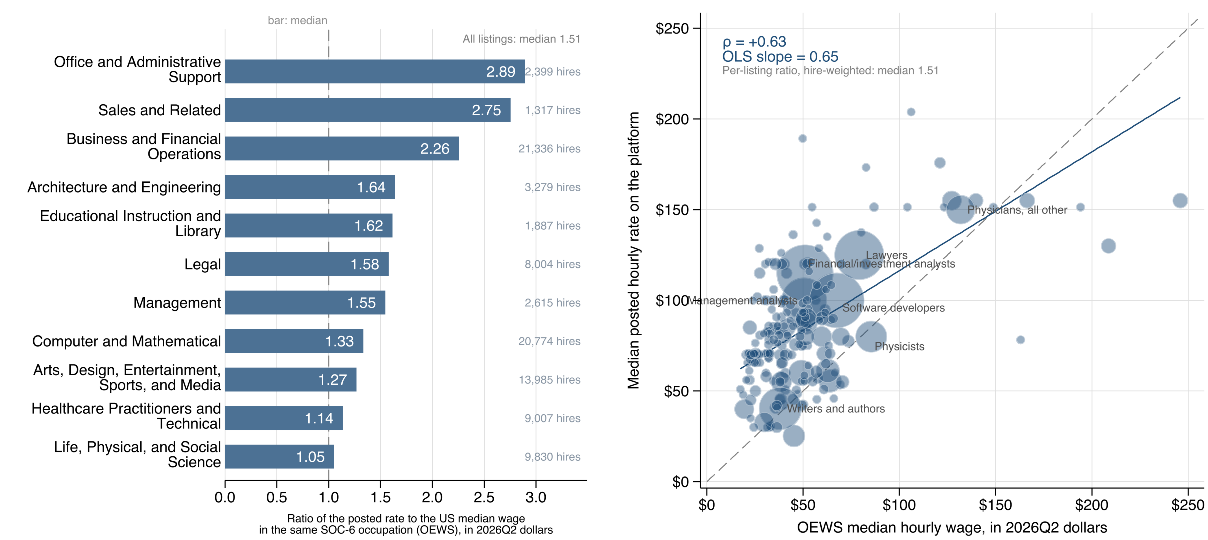

A. Median pay ratio by SOC-2 (midpoint over OEWS median), weighted by hires

B. Median posted midpoint against the US median wage, by SOC-6, weighted by hires

Figure 15: Robustness: the floor of the posted range

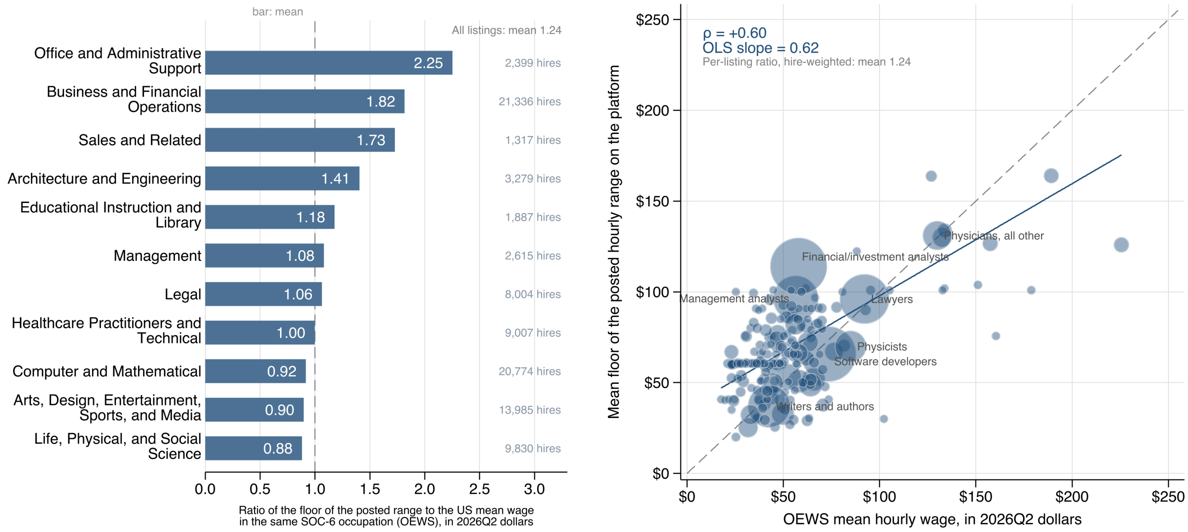

A. Mean pay ratio by SOC-2 (range floor over OEWS mean), weighted by hires

B. Mean posted floor against the US mean wage, by SOC-6, weighted by hires

Figure 16: Robustness: every listing weighted equally

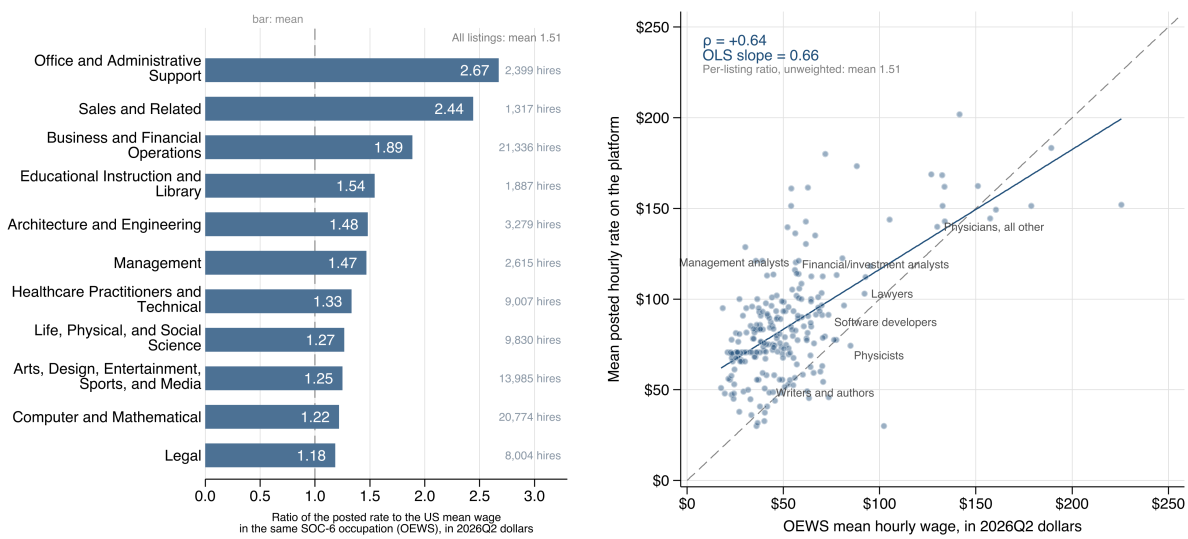

A. Mean pay ratio by SOC-2 (midpoint over OEWS mean), unweighted

B. Mean posted midpoint against the US mean wage, by SOC-6, unweighted

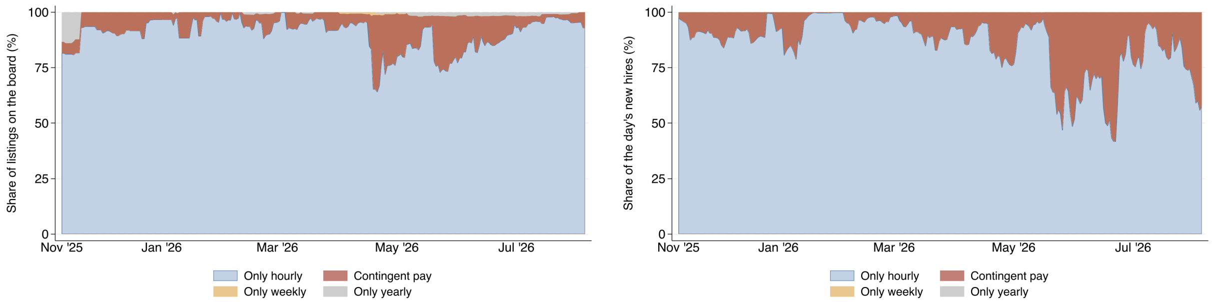

A. Share of the listings on the board B. Share of the day's new hires, 7-day window

Figure 17: Composition by pay basis, weighting the board two ways: (a) each listing counts 1, so the bands are the composition of the board; (b) each listing counts by its new hires that day, so the bands are the shares of hiring, over a centered 7-day window. The three time-rate bands are plain time rates only: a listing posting an hourly, weekly or yearly rate and nothing else. *Contingent pay* is everything else — a fixed payment for the job (one-time or per-task) *plus* any listing whose description that day offers a cash bonus on top of the rate or promises a rate, role or pay-format escalation on merit.

Figure 18: Correlation of market size and prices with AI-use measures

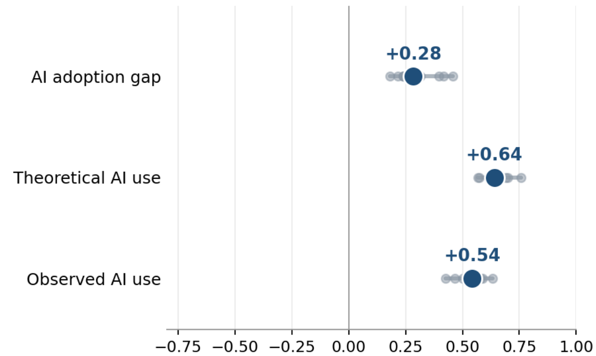

A. Total hires (log)

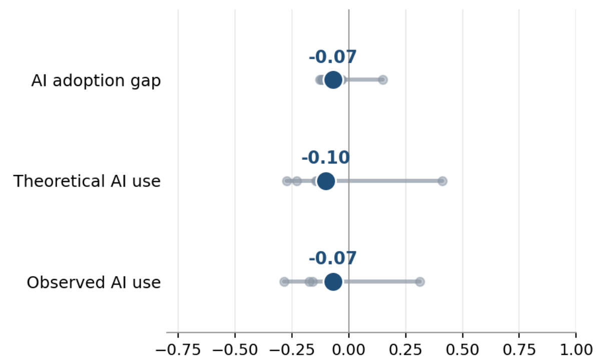

B. Median posted price (log)

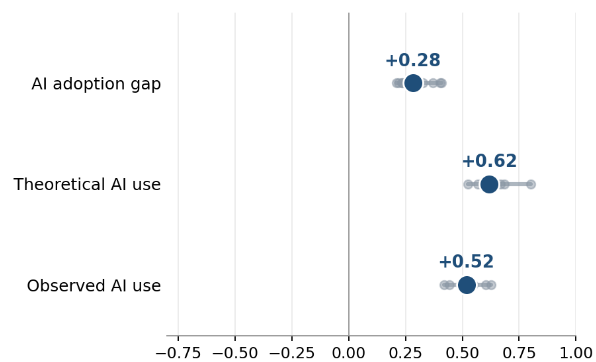

C. Estimated spend (log)

Notes: Correlation between three market outcomes and the three AI-use measures of Figure 6, across the same twenty groups, likewise weighted by hires, and over the same November 11, 2025 to August 5, 2026 panel, with leave-one-out ranges constructed identically. Total hires and estimated spend are posting-level sums over each group's listings (Data Description), in logs; the median price is the median posted midpoint across the group's hourly listings, in logs. Market size loads on theoretical exposure: hires correlate +0.64 with theoretical AI use against +0.28 with the adoption gap, and estimated spend +0.62 against +0.28 — groups where AI could absorb more work bought more expert labor and spent more. Posted prices are unrelated to all three measures (−0.10 to −0.07), with leave-one-out ranges that cross zero in every case.

Figure 19: Timing of expert hiring and the AI adoption gap, groups above 1,000 hires

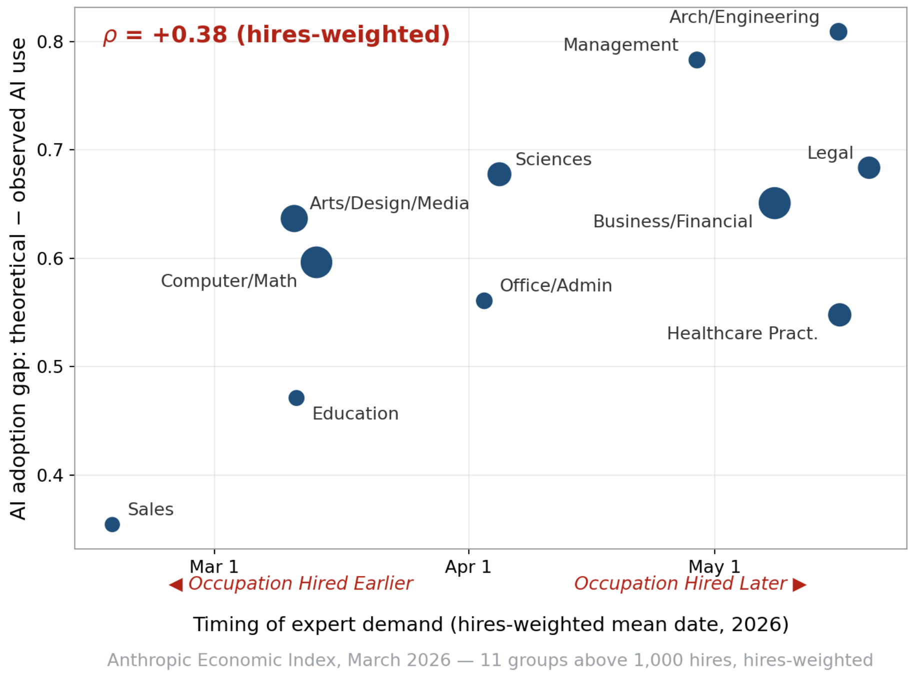

Notes: Figure 5 restricted to the eleven major groups with at least 1,000 hires over the panel, the selection rule used elsewhere in the paper; axes, weighting and construction are otherwise identical. The nine groups this drops account for 1,243 hires, 1.3 percent of the total, so the weighted correlation barely moves, from +0.41 to +0.38.