August 2026

Hiding in Front of Us: AI Impact on Jobs

Liminal Capital.
Ben Verschuere - Co-CIO
ben.verschuere@liminal-capital.com

Angus Cameron - CIO
angus.cameron@liminal-capital.com

The argument about whether AI is destroying or creating jobs is not a disagreement about the data. We argue the focus has been on the wrong object and instead provide an alternative framework with answers.

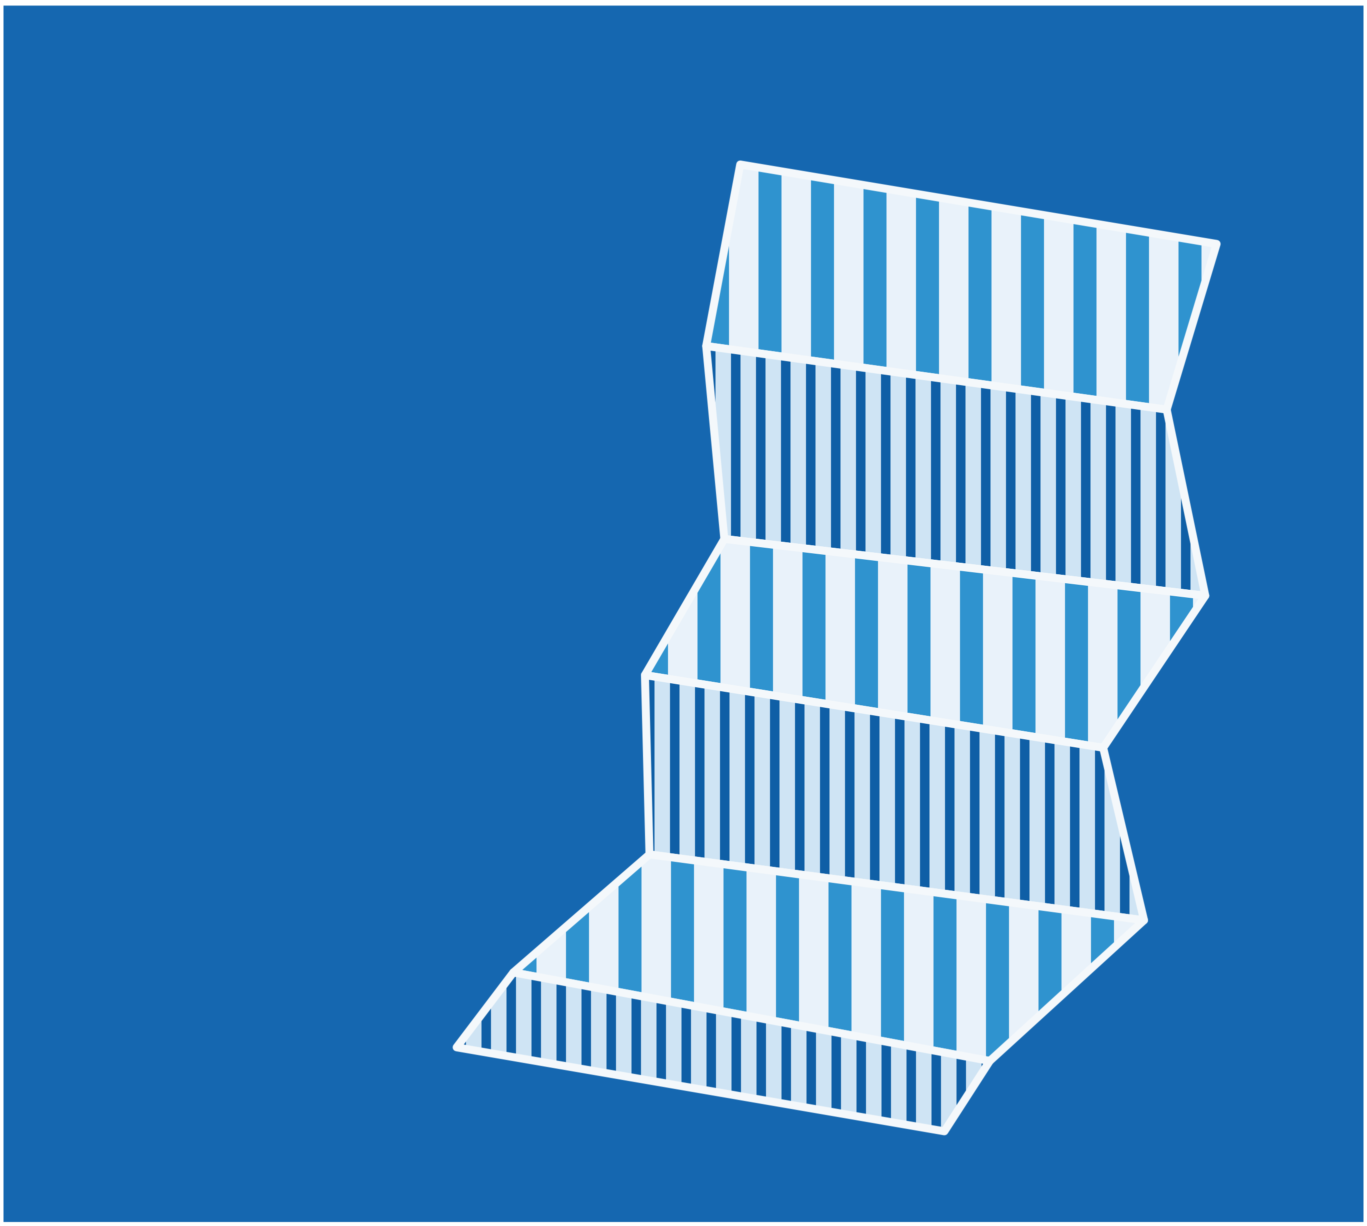

EXECUTIVE SUMMARY

AI automates *tasks*, not *jobs*, but so far almost every published AI exposure index is built at the *job* level, which hides AI's two opposite effects. On some tasks AI acts as a substitute for the human, on others it acts as a complement. Both happen inside the same job, so a single job-level number records the net of the two and produces ambiguous results. That is why published estimates of AI's employment impact can disagree not only in size but in direction.

We estimate that AI is a substitute for 32% of American work, a complement for 15%, and has no effect on the remaining 53%. We graded all 17,536 task statements in the government's occupational database, then aggregated them back up to the job. Instead of one exposure score per job, every occupation now carries three numbers: substitution, complement, and no effect (ie the job is physical and AI cannot touch it).

Work AI can substitute has grown 3.3pp a year more slowly than work AI does not touch, while work AI complements has grown 3.3pp a year faster: a net drag of about 800,000 jobs a year. Neither effect existed before 2021, and neither is COVID unwinding. In headcount that is roughly 1.5 million jobs a year removed through AI substitution against 0.7 million added through complementarity (measured against the same work being out of AI's reach).

The effect works through less hiring rather than firing, and it lands on the young. Job starts for workers aged 22 to 25 in the most substitutable occupations are down by a third, some 227,000 entry-level positions a year that no longer open, with no matching rise in firings. Within substitutable work the effect is also propagating up the pay scale, having turned significant in middle-paid jobs over the last year.

RIGHT DEBATE, WRONG OBJECT

At its core, **AI automates *tasks rather than occupations***, and this distinction may be the single largest reason published estimates of AI's employment impact vary so widely.

Consider how much AI speeds up a *job* such as a lawyer. The question cannot be answered without knowing which *tasks* a lawyer performs. AI drafts a brief in minutes and does nothing at all for a court appearance. The technology does not act directly on the job. It acts first on individual tasks, and a job is a bundle of heterogeneous ones.

The academic literature has treated *tasks* as the correct unit since the early 2000s, but the measurement practice built on top of it went the other way: nearly every standard AI exposure index scores at the *occupation* level. The reason is practical. The US economy has roughly 800 occupations against 17,536 distinct tasks, so grading occupations was tractable and grading every task by hand was not. Large language models have since removed that constraint, and we use them for the grading described in the next section.

Measuring AI exposure at the job level destroys exactly the information that matters.[1] Take two occupations the most established AI exposure index (the AI Occupational Exposure, or AIOE) scores nearly identically: medical secretaries & administrative assistants at +1.218, secondary school teachers at +1.191 (0.027 apart on an index spanning several points). For the medical secretaries, 67% of the work is what AI can do on its own, acting as a substitute for the human, and 7% needs a human present, with 26% the physical remainder. For the teachers the ratio is inverted: 22% substitutable, 36% acting as a complement and 42% physical. Using the most prevalent AI exposure index both jobs look identical, but inside them the task composition is completely different (and between them they employ nearly two million Americans).

1. Some of our previous work missed this as well by using an aggregate measure of AI exposure (*C**). While it provides a decent first-level approximation at the economy-wide level, it is too crude to detect AI's effect at the job level; which the modelling presented here solves for.

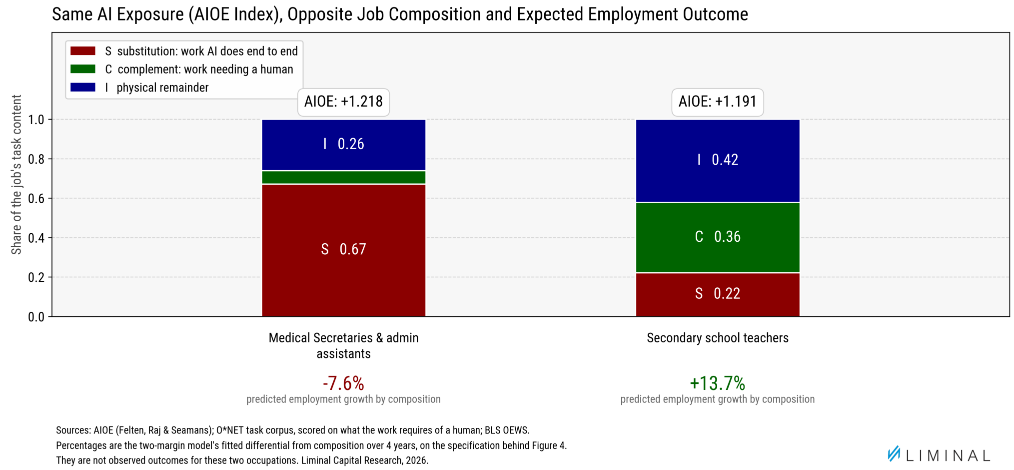

Figure 1. Two occupations the published index cannot separate, broken into substitution, complement and physical work. The percentages below the bars are what our model implies from each job's makeup, not what was observed for these two occupations.

The split between substitution and complement (and physical) is what decides the direction of AI's effect on a job. Where AI substitutes, it does the task instead of the worker and fewer workers are needed. Where AI complements, it produces part of what the worker delivers, cuts the cost, and more workers can be needed. As we have seen in the example above one general AI exposure number is too crude to see the difference. Furthermore it can only rank jobs by how badly they are hit, but it can never name the ones that are helped (or vice versa).

Instead, two measures (the complement share *C* and the substitutable share *S*) carrying opposite-signed coefficients can. And because the complement and substitution forces offset, a single number understates the damage where AI substitutes and overstates it where AI complements. This error cancels in the aggregate while being potentially wrong about every occupation inside it.

WHAT THE WORK IS ACTUALLY MADE OF

Instead of assigning *one* AI exposure metric to each job, our framework scores each of the 17,536 task statements in the government's occupational database, sorting every task by asking a simple question: *what does this task require of a human?*

What the task requires of a human	Economic interpretation
Nothing. The output is text, code or analysis, and AI produces it.	**Substitution, *S***
A human present in real time, or a licensed human to sign.	**Complement, *C***
It is physical work. AI cannot perform the task at all.	**Inert, *I***

The question asks what the work requires, not what today's models happen to be able to do, so the answer does not move when the models improve.[2]

With every task scored as *S*, *C* or *I*, each job (ie a bundle of tasks) receives its own values, between 0 and 1, for *S*, *C* and *I*. Each occupation then carries three measures rather than one crude exposure number: its substitution share *S*, its complement share *C*, and the physical remainder *I* that AI does not touch (this is what Figure 1 showed earlier for two different jobs).

Having established the scoring, this lets us estimate that across 143 million US jobs the economy splits into 45 million jobs' worth of substitutable work (32%), 21 million of complementary work (15%) and 77 million of physical work (53%) that AI cannot touch. This provides an aggregate mapping, but since each job is a bundle of different types of work we can build a more precise topography, which is what Figure 2 provides.

2. The scores are produced by grading every O*NET task statement against a fixed structural rubric. Because the criterion asks what the work requires rather than what a model can currently do, it is stable across raters: six large language models from four different developers agree unanimously on 82% of tasks.

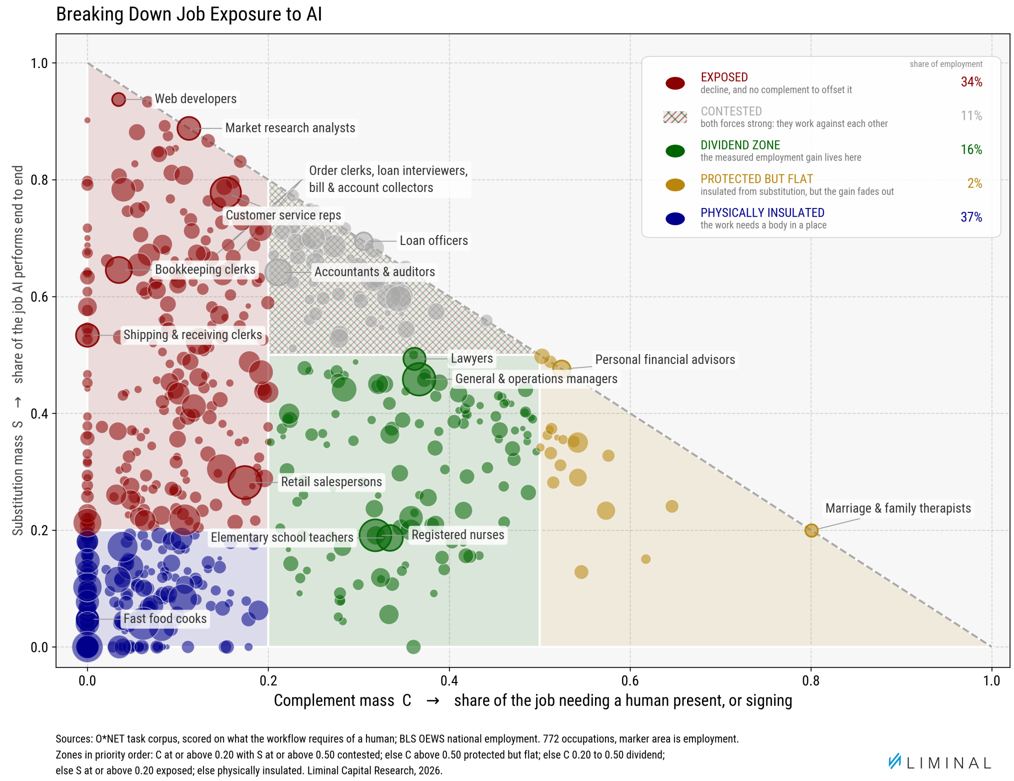

Figure 2: Mapping every US occupation by its substitution and complement shares

Drilling deeper, Figure 2 places every occupation by its own *S* and *C* values. Each bubble is one occupation, sized by how many Americans work in it. Its height is *S*, the share of the job AI performs end to end (at *S* = 1 AI can do the whole job; web developers sit close to that). Its horizontal position is *C*, the share needing a human present or signing, where AI can act as a complement (because the shares sum to one, no job can sit above the dashed diagonal). Finally, a point sitting low and left, where both *C* and *S* are close to zero, is by definition a physical job AI cannot do (since S + C + I = 1, if *S* and *C* are both small then by definition *I* is close to 1).

Having identified the topography it is worth looking at the partitions inside it.

As mentioned, dots in the lower-left corner are by definition physical and are the ones AI does not touch. This is the **blue region**, around 37% of employment, which AI does not impact because the work is physical; fast food cooks, for example.

Moving up, the **red region** is where AI has the greatest impact: jobs with high AI-substitutable content and little complementary work to absorb it, around 34% of employment; including bookkeeping clerks, market research analysts and web developers. These are the jobs AI can directly replace, and the most likely to have felt an impact since the launch of ChatGPT.

It is worth noting that the top of that red strip is the sharpest reading in the figure. One in six American workers is in a job that is majority automatable with little human-contact work to absorb the freed capacity. That is 22.3 million people, and it is concentrated: the ten largest such occupations hold 54% of all at-risk employment. They are customer service representatives (2.7 million, 78% automatable), secretaries and administrative assistants, bookkeeping and accounting clerks, business operations specialists, receptionists, management analysts, market research analysts, shipping and receiving clerks, medical secretaries and executive assistants.

At the centre of the triangle lies the **green region**, the safer zone, representing around 16% of employment. These occupations have a mix of mild AI substitution and complementarity. Jobs such as nurses and school teachers sit here: AI can help with part of the work but the job still requires a human present. Using AI improves productivity (as well as reducing the cost) *without* removing the human the job needs, which then raises demand for the output the job produces and therefore headcount. This is in essence the radiologist example Jensen Huang likes to use: as they become more efficient thanks to AI, it drives up demand and drives up employment for radiologists, who are still needed to interact with patients. The green region is where AI can drive employment upward.

Moving to the bottom right, the **gold region** holds work where AI acts purely as a complement and substitutes for nothing, such as marriage therapists. A human is needed throughout, so AI does not move employment much: the cost of an hour of therapy barely changes with or without AI. Those jobs represent a small fraction of the labour force, around 2% of employment.

Lastly, along the upper part of the map lies the **grey region**, a contested zone of about 11% of employment, where the positive effect of AI as a complement works against its effect as a substitute, so the outcome is genuinely contested. These are jobs such as accountants and loan officers.

That map is also why AI's impact on jobs looks calm in the aggregate. The blue block and the gold band together hold 39% of American employment, where AI has little impact: nursing assistants on their feet, electricians, drivers, cooks at one end, therapists and advisers at the other. Nothing AI does registers in that part of the chart. The action is in the remaining three zones, and there, in terms of employment share, the red outweighs the green by roughly two to one, 34% against 16%.

Before using this framework to analyse employment, it is worth a brief detour into a driver that sits between the task measurement and the job outcome which is the demand for what these jobs produce. Automating part of a job cuts the labour needed, which lowers price and then potentially raises the demand for the output produced by those jobs. Whether employment falls depends on which dominates. We do not measure that demand response (ie the elasticity) and do not need to: what we measure is what actually happened to employment, which already contains it. But it is one reason AI's effect is not uniform. Where demand rises enough when the price falls, an occupation can be highly substitutable and still grow (the next section shows the one place that is currently happening, the technology occupations that build AI itself, which on composition alone should have been shrinking but are instead growing due to the AI boom).

BREAKING DOWN AI IMPACT ON JOBS

With both *C* and *S* measures in hand for every occupation, we can look at whether employment growth since ChatGPT depends on them. The window runs from 2021, the last full year before the launch, to 2025, the most recent employment data at the occupation level.

AI-substitutable employment declined; AI-complementary employment increased. Substitutable work grew 3.3 percentage points a year more slowly than work AI does not touch. Complementary work grew 3.3 points a year faster. Both results are statistically significant.[3] The fact the two numbers come out the same size is a coincidence, not something the model imposes (both are also measured relative to the rest of the economy, so neither is affected by how the broader economy moved).

At the economy level, combining the two effects, **we estimate that AI has so far removed about two jobs for every one it created, a net loss of 800,000 a year**. Substitution removes roughly 1.5 million jobs a year; complementarity adds roughly 0.7 million. The two-to-one ratio is not because substitution bites harder than complementarity helps. The two happen to run at the same speed, instead it is because the economy holds twice as much substitutable work as complementary work: 45 million jobs' worth against 21 million. The ratio is a fact about what the economy is made of, not about the strength of either effect.

Figure 3 traces both paths from 2015, each built up period by period. Before 2021 neither can be told apart from the rest of the economy, which is what the dotted segments mean. After 2021 both are significant, and they finish 18 points apart. Each line also passes its own 2019 level and continues past it, which is what rules out a COVID unwinding.

3. Substitution p = 0.003 under standard errors, 0.067 under the more demanding test the underlying paper adopts. Complementarity p = 0.031, essentially unchanged at 0.029 under the same. The paper grades the complement result robust and the substitution result consistent with AI rather than established.

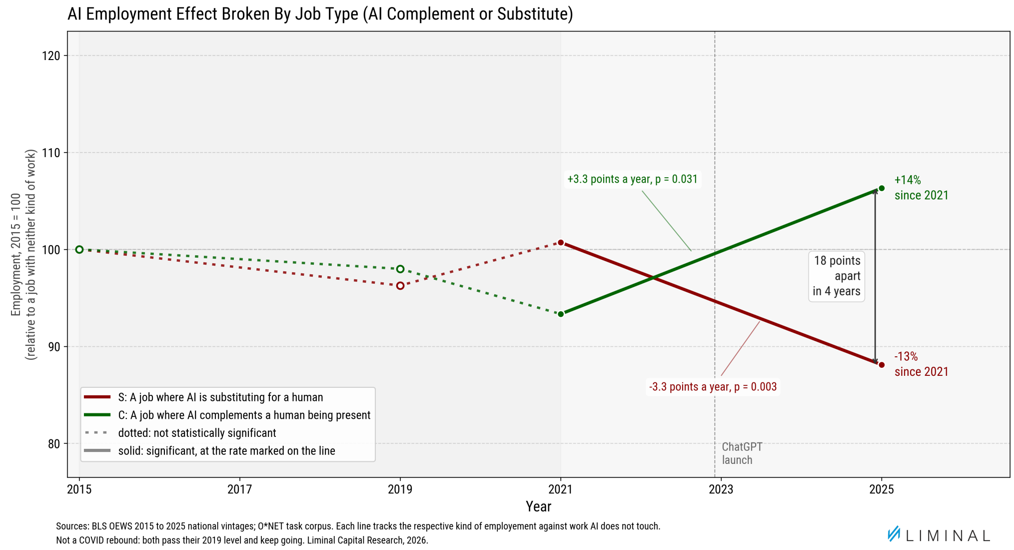

Figure 3: Relative employment paths for substitutable and complementary work, 2015 to 2025

It is worth pointing out that the substitution and complement effects scale differently. Substitution is proportional: the more substitutable a job is, the more its employment has fallen, shown by the downward-sloping line in Figure 4. There is nonetheless one exception, at the very top. Eight computer occupations, the ones that build and deploy the technology, are growing despite being among the most substitutable in the economy. This is the demand channel from the previous section working in the other direction: demand for AI's own supply chain has risen faster than AI has displaced the work, so high substitutability has not translated into fewer jobs.

The complementarity dynamic is not linear but instead an inverted U, with no effect at either end and the gain concentrated in the middle. It is worth highlighting this dynamic. Theoretically, at one extreme where C = 0 there shouldn't be any AI impact at all given the job is highly physical (this is the blue region in Figure 2 where AI has no impact). While at C = 1 the job is delivered entirely by a human, the marriage therapist case (the gold region in Figure 2) where again AI has also no impact. Interestingly, as Figure 4 shows our estimates recover in practice those null effects at both boundaries.

Figure 4 furthermore puts numbers on both the substitution and the complement impact. As shown on the left-hand side a job that is 80% automatable carries about a 16% employment shortfall over four years, while on the right-hand side the complement arch peaks near +20% at C = 0.31 and is back to nothing by C = 0.61. These relationships map back onto the Figure 2 topography: jobs where S is high and C is low are the most negatively affected (the red region), while jobs where S is relatively low and C sits around 0.3 are positively affected by AI; this is the radiologist example (the green region).

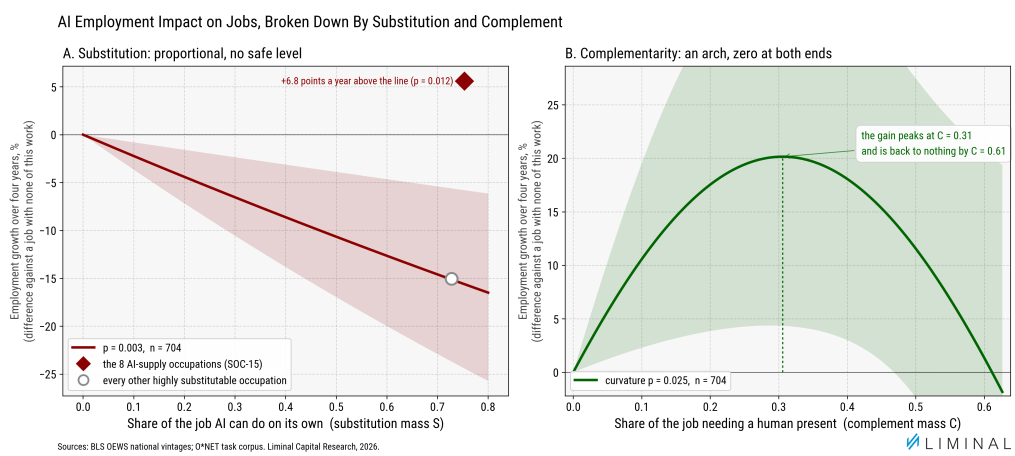

Figure 4: Employment impact depending on how much substitution and complement a job contains

Going back to the single AI exposure metrics one might wonder how our *C* and *S* measures interact with it. We find that when we combine our *C* and *S* breakdown with the typical single measures of AI exposure, those single measures lose their significance. One-number measures are simply too crude: whatever they were picking up, *S* and *C* pick it up better. As mentioned earlier, those single measures also cannot account for AI both adding and removing jobs at the same time (as one number cannot carry two signs).

Finally, a natural objection to our analysis might be that we have shown timing, not causality. It could be argued that the two effects arrive alongside AI, which is not the same as being caused by it, and they could be picking up something else. To test for that we controlled for different effects: a post-COVID reversion, the 2022 to 2023 interest-rate cycle that raised the cost of capital, the long-running decline of clerical work, and the shift to remote work. None of these removes the significance of the *S* and *C* effects. We also ran the identical test entirely inside 2015 to 2019, before AI, and neither effect appears. Of the explanations we tested, AI is the one left standing.

THE HIRING GATE: WHO IT HITS, AND WHERE IT GOES NEXT

Having established how AI affects jobs differently depending on the tasks that compose them, we turn to the employment dynamic this creates, who it affects, and where it goes next.

The AI effect is a hiring gate rather than a firing gate, and it currently falls on younger and lower-paid workers. Job starts for workers aged 22 to 25 fell by a third in the most AI-substitutable occupations, about 227,000 entry-level positions a year that no longer open. Hiring of the same cohort into work AI *cannot* substitute did not decline over the same period, so this is not a general freeze on young hiring but instead one directly focused on substitutable jobs (we also find no significant rise in firings). This hiring gate is a similar effect to the one Brynjolfsson, Chandar and Chen (2025) find, and to what we documented in Verschuere and Cameron (2026a).

This gating mechanism on the young is a puzzle. What we know is that it is not significant for older workers and absent for complement jobs at every age, and it is not driven by occupation mix within the younger cohort nor by the sectors those workers sit in. The same mechanism is also documented in payroll, resume, vacancy and register data across the United States, the United Kingdom and Sweden, and no published work has identified why it falls on the young. A potential explanation of this "quiet hiring" could be that it is what a cautious response to a new technology looks like. A tool arrives that can do part of the work. Firms adopt it, up to a point, and adjust by hiring less of the cohort which has the highest entrant to employment ratio (ie the young). None of this requires anyone to be let go. There is no separation, no notice and no announcement, and the adjustment happens in vacancies that are never posted.

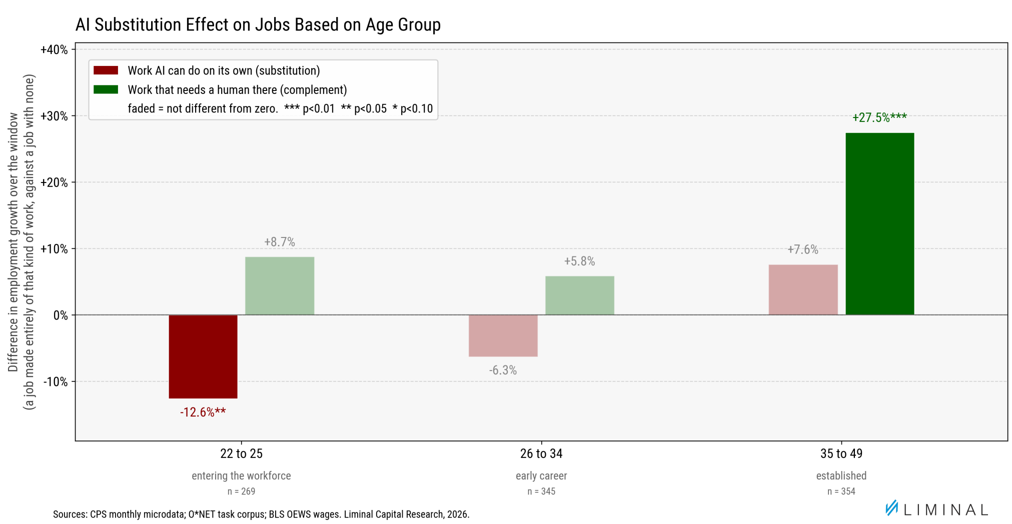

Figure 5: The employment impact from substitution and complement effects by age band

Further exploring the dynamic of this quiet hiring also means the age profile of those affected should not stay where it is. A worker not hired at 23 in 2023 is a missing 26 year old today. We see some indicative information about that from the 26-34 year cohort, which has been shrinking, though not significantly, while we do not see any significant impact for older employees in these highly substitutable jobs. On the flip side, for jobs where AI can act as a complement we actually see significant employment growth for older workers.

The negative AI effect on substitutable jobs is climbing the pay ladder. As shown in Figure 6, another dynamic we can observe is the dynamic across pay cohorts (instead of age). We find a similar effect where employment at the bottom of the pay distribution has been declining (−7% per year) since ChatGPT as shown in Figure 6. What has started to change is the middle-paid jobs in exposed occupations, where the decline in jobs has started to accelerate (from −2.6% to −4.5%), while the highest paid cohort, which was still growing from 2021 to 2024, has now flipped sign, though not yet at a significant level. Taken together, this shows the effect of AI working its way up the pay scale for substitutable jobs.

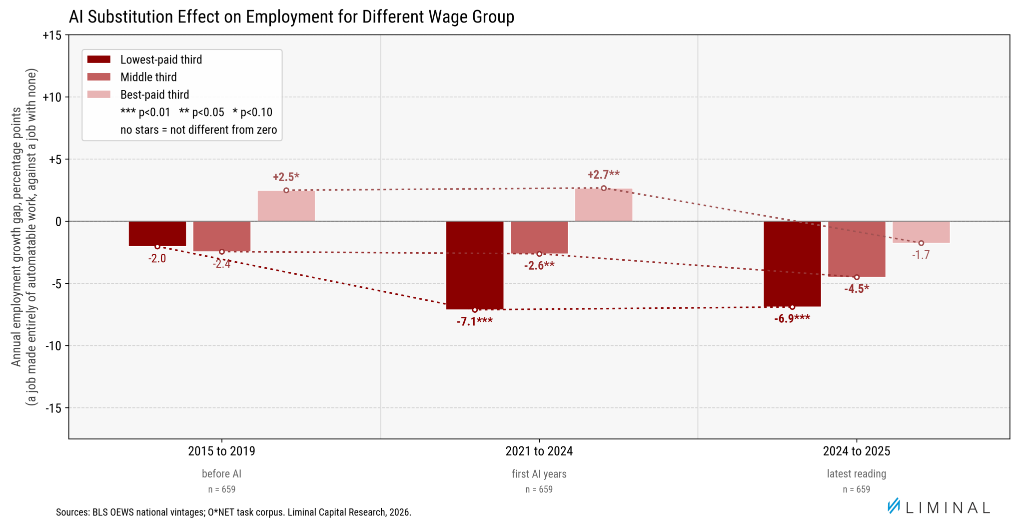

Figure 6: The employment impact from the substitution effect by wage cohort

What we have shown so far is that lower-wage and younger workers in substitutable work have been hit hardest. It is natural to wonder whether such a negative effect on employment also transfers to wages. In short, we have not found one. That may be partly because two forces run against each other. A lower hiring rate in those cohorts mechanically drives the average wage higher, since new entrants typically get paid less and fewer of them are entering, while weaker demand for substitutable labour pushes wages down. The two offset, which makes any wage effect hard to identify in our specification.

Overall we could describe the AI effect on jobs so far as a "quiet hiring" effect, disproportionately affecting higher-turnover, lower-paying and entry-level jobs. The point at which this stops being quiet is a recession. As we have seen, lower paid jobs have already been hit by the hiring gate, while higher-wage occupations may have much more room for deeper cuts, compounded by the higher salaries those jobs carry. Those may be the jobs most affected when the next downturn happens.

POLICY IMPACT

If AI were producing layoffs, the existing unemployment benefit apparatus would work as designed. A worker who is never hired triggers none of it. No separation, no claim, no notice, no caseload. A hiring-margin transition is one the safety net neither cushions nor counts, and political support cannot easily be assembled for a response to something that generates no statistics.

Timing compounds the problem rather than deferring it. The research on entering the labour market under poor conditions is unusually consistent: affected cohorts carry earnings losses for decades, missing a favourable entry window does more damage than enduring an unfavourable one, and the least advantaged entrants are harmed most.

A constraint on entry does not postpone its cost until conditions improve. It compounds cohort by cohort for as long as it persists.

There is nonetheless a silver lining. As we have documented, a meaningful number of jobs are benefiting from AI: the green dividend zone of Figure 2, roughly 16% of employment. The potential issue is a skills mismatch: entrants to the job market have been responding to a pre-AI signal, pointing them toward jobs that were reliable entry points into professional work for two decades, so the cohort now arriving was steered toward them by a signal that has since inverted. What matters is how quickly the new signal propagates back into training and course choice. Our framework can not only identify where jobs are accruing thanks to AI, but also provides a mechanism to track the phenomenon.

CONCLUSION

The debate over AI's effect on employment has not lacked arguments or data. It lacked the right object of measurement. When focusing on the right object, the task, we can observe a significant AI employment impact from which we can draw meaningful take-aways.

The first is the presence of two opposite effects. AI substitutes for some work and complements other work, and a single number records only the difference between them. Measured separately, the two run at almost exactly the same speed, 3.3 percentage points a year in opposite directions. The economy-level result is negative only because there is twice as much substitutable work as complementary work: 45 million jobs' worth against 21 million.

The second only becomes visible once the substitutable work can be isolated, and it is the dynamic inside it. Firm are hiring less, and almost all of the measured effect sits in entry hiring into substitutable work, down by a third. The effect is nonetheless propagating up the pay scale within substitutable work, and that is the direction to watch.

It is worth noting that the coefficients we estimate are not set in stone, and it remains to be seen how the substitution and complementarity gradients change over time, and how the impact propagates through the pay scale; especially given that four years into a major technological development is a short window.

Any technology that acts on *tasks* will be mismeasured by *occupation*-level averages whenever two effects inside it oppose. The framework we provided here measures those separately, and thus maps where the AI employment headwind and the tailwind fall. What could hide behind a single AI exposure, a decomposition between AI-substitution and complement can now see.

SOURCE AND REFERENCE

This note is based on the underlying paper: Verschuere and Cameron, *Hiding in the Mean: The Two Margins of AI's Employment Effect* (July 2026), SSRN: link.

Verschuere, B., & Cameron, A. (2026a). *AI and the Economy: The Four Rates Which Define The Future.* Liminal Capital working paper, April 2026, current version May 2026.

Brynjolfsson, E., Chandar, B., & Chen, R. (2025). *Canaries in the Coal Mine? Six Facts about the Recent Employment Effects of Artificial Intelligence.* Stanford Digital Economy Lab working paper, November 2025.

Hosseini Maasoum, S., & Lichtinger, G. (2025). *Generative AI as Seniority-Biased Technological Change: Evidence from US Resume and Job Posting Data.* SSRN working paper 5425555.

Lodefalk, M., Lothman, A., Koch, M., & Engberg, E. (2026). *Generative AI and the Age Gradient in Hiring.* Orebro University School of Business working paper 2/2026.

Westby, S., Modestino, A. S., & Cheng, Z. (2026). *Rising Experience Requirements and the Entry-Level Software Market.* IZA Discussion Paper 18723.

Humlum, A., & Vestergaard, E. (2025). *Still Waters, Rapid Currents: The Labor Market Effects of Generative AI in Denmark*. NBER Working Paper 33777.

Eloundou, T., Manning, S., Mishkin, P., & Rock, D. (2024). GPTs are GPTs: Labor market impact potential of LLMs. *Science*, 384(6702), 1306–1308.

Felten, E., Raj, M., & Seamans, R. (2021). Occupational, industry, and geographic exposure to artificial intelligence: A novel dataset and its potential uses. *Strategic Management Journal*, 42(12), 2195–2217.

Data, all public: O*NET task database; BLS Occupational Employment and Wage Statistics (2015 to 2025 vintages); BLS Current Population Survey microdata; BEA real value added by industry; the AIOE index (Felten, Raj and Seamans) and the task-exposure scores of Eloundou et al., referred to in the text as the published exposure measures. The payroll corroboration is Brynjolfsson, Chandar and Chen (2025). The sector-level and JOLTS evidence on the hiring gate is Verschuere and Cameron (2026a).