

Capital Reads Faster Than Labour

Three AI stories, one feedback loop, and a market that has stopped telling firms who is good.

CAPITAL READS FASTER THAN LABOUR

Private Equity Is Deploying AI Before Productivity Is Proven

Why deployment capital now moves faster than organisational readiness

The hidden feedback loop between AI hype, asset protection, and workforce disruption

Inside the new race between capital allocation and operational reality

Inside the new race between capital allocation and operational reality

AI capability is improving slowly. Market pressure is not.

The firms setting the pace are not CIOs. They are capital allocators.

The most-discussed AI story of the past month is not the Anthropic–Wall Street joint venture, the Anthropic–xAI compute deal, or the [Noema essay](#) arguing that liberal arts graduates will inherit the AI age. It is the **invisible relationship between the three of them.**

Read in isolation, each is a familiar headline. Read together, they describe one feedback loop, and that loop is not really about job destruction. It is about something quieter and, in the long run, more damaging:

Before AI takes anyone's job, it takes the market's ability to tell who is good at the job.

That sentence is the load-bearing claim of this piece. It explains why corporate AI deployment is racing ahead of corporate AI capability. It explains why white-collar work is being reshaped before any AI lab can pass a serious benchmark. And it explains why the labour market itself, the application, the CV, the LinkedIn post, the screening call, has become a hall of mirrors that nobody, on either side, can navigate cleanly.

Let me show you the loop.

Story 1, Capital is moving ahead of capability

On 5 May 2026, the Financial Times reported that **Anthropic, Blackstone and Goldman Sachs** had launched a joint venture to embed Anthropic engineers inside the private equity firms' portfolio companies, the mid-sized US businesses where AI disruption is expected to bite first. The Perplexity coverage of the same week put the figure at \$1.5 billion and added Hellman & Friedman to the list. OpenAI announced a parallel \$10 billion vehicle, *The Deployment Company*, on the same day.

That timing is the story. Two rivals, two associations, one synchronised push to **buy distribution into mid-market firms whose value is at risk**. The FT's Ryan McMurrow framed it bluntly: PE firms fear that their portfolio companies will be "oozed out" by AI before they can exit them.

Now read that against the Anthropic-xAI announcement from the same period: Anthropic, the company building these tools, is **so capacity-constrained that it is sharing GPUs with a rival**. The deal gives Anthropic access to xAI's Colossus supercomputer in Memphis, roughly 200,000 GPUs at nearly 2GW, to relieve the throttling that has frustrated Claude Pro and Max subscribers for months. This is on top of a \$100 billion AWS commitment, a multi-gigawatt agreement with Google and Broadcom, and chip talks with the UK startup Fractile.

This is unambiguous. **Demand is outrunning supply at every layer**. Compute is being rationed. Capital is being routed into deployment vehicles. EPAM has trained 20,000 of its own staff on Claude through Anthropic's partner network. None of this waits for productivity reports.

If you are a corporate executive, this is the operative fact. The pace of AI adoption inside your industry is no longer being set by your CIO's pilot programme. **It is being set by capital allocators who have decided that not deploying is more dangerous than deploying**. That is a different speed, and it is very different politics.

Story 2, The work has redistributed upward, even though AI is failing the work

This is the contradiction.

In January 2026, the training-data company **Mercor** ran a benchmark study on whether AI agents could perform real white-collar tasks drawn from consulting, investment banking and law. Every major AI lab failed. Even the best models scored under 25%. TechCrunch reported the study under the headline "*Are AI agents ready for the workplace? A new benchmark raises doubts.*" A King's College London study from October 2025 reached a similar conclusion from the labour-market side: software engineering and

management consultancy were the most exposed to wage and employment effects from AI, but mass displacement had not arrived.

Joe Procopio's diagnosis in *Entrepreneurship Handbook* is the severest I have read: "every AI lab is getting a failing grade" at replacing knowledge workers. His explanation is simple and correct. **Knowledge work is not keyboard-pushing. It is the application of context, judgement and improvisation to messy environments**, which AI cannot access because that context lives inside human heads, not in structured data online.

And yet.

Anthropic's CEO Dario Amodei has confirmed, repeatedly, that **over 90% of the code now powering Claude is written by Claude itself**. Boris Cherny, the creator of Claude Code, said in December that 100% of his own contributions that month came from the AI. Anthropic shipped its consumer agent product *Cowork* in roughly ten days using Claude Code as the developer.

So which is it? Is AI failing at white-collar work, or is it eating its own engineering function?

Both. The disruption is not a clean replacement. It is a **vertical redistribution**.

Amodei was explicit on this point in his Dreamforce conversation with Marc Benioff, reported by Business Insider: "If Claude is writing 90% of the code, what that means, usually, is you need just as many software engineers. They can focus on the 10% that's editing the code or writing the 10% that's the hardest, or supervising a group of AI models."

Read that as an executive. The juniors do not get hired. The seniors do not get fired. The work shifts to the 10% that is hardest, supervision, edge-cases, the exact zone where AI fails the Mercor benchmark. The labour-market consequence is not destruction. It is a **collapse of the bottom rung of the ladder** while the top rung becomes more demanding.

The Guardian's February piece "*The big AI job swap, why white-collar workers are ditching their careers*" documented the human side: writers, editors, lawyers, occupational health professionals retraining as therapists, bakers, electricians. Carl Benedikt Frey at the Oxford Internet Institute, the same economist who in 2013 warned that 47% of US jobs were at risk from automation, has now revised his message: "*We're beginning to see some studies suggesting more of an impact on entry-level work.*"

Nils Gilman's *Noema* essay, published 9 April 2026, names the danger that this redistribution conceals. Entry-level jobs are not just where juniors earn their salary. They are the apprenticeship through which juniors observe seniors and acquire the tacit judgement that machines cannot replicate. Strip that rung, Gilman argues, and you are "*eating your own seed corn.*"

A firm can run on senior judgement for five years. Not for fifteen.

Story 3, Both sides of the labour market are now AI-mediated

The third story is the one that turns the first two into a closed loop.

A February 2026 study by Publicate, reported in *TechRadar*, found that **73.8% of LinkedIn posts in finance and 57.8% in technology now show signs of AI authorship**. Mid-level professionals lead adoption (54.4%), the exact population fighting hardest to prove they belong on the senior side of the new redistribution. Senior executives follow at 53.6%. Entry-level staff at 52.2%, compensating for thin experience with polished output.

That is the supply side of the labour market presenting itself through AI. Now the demand side.

The CEO of Greenhouse, Daniel Chet, told an FT podcast that recruiters now process roughly **3.5 times the application volume they handled in 2022**, while global hiring sits 20% below pre-pandemic levels. His firm processed 65 million job applications last quarter alone. The volume is not an accident: candidates use AI to apply more, and faster. Employers respond with AI screening tools. New York's Local Law 144 and California's hiring AI rules are attempting to regulate the symptoms.

A separate FT analysis from 20 February 2026 reported the consequence of this from the graduate side: Hannah, a King's College graduate in war studies and international relations, applies into "the void" and is ghosted dozens of times. She is not alone. Greenhouse data confirms what graduates already feel: applications must be tailored to algorithms first, humans second. AI-on-AI screening is now the dominant interaction in the early-career labour market.

Then, the darker corner. Chet again, on the same podcast: deepfake interviews, North Korean impersonation rings infiltrating remote IT contracts, candidates piping ChatGPT through their headphones during live calls. The infrastructure that allowed remote hiring to scale is the same infrastructure being weaponised against it.

What you have, at the level of mechanism, is a labour market in which **both sides are using AI to read the other side, and neither trusts what it sees.**

The loop

Now stack the three stories.

1. **Capital deploys AI ahead of capability.** PE firms cannot afford to wait, so the deployment happens whether or not the technology delivers.
2. **The work redistributes upward.** Juniors are not hired; the senior layer absorbs the failures of the AI as supervision and edge-case work.
3. **Workers signal AI fluency to compete for what is left.** The most legible proxy for "ready for senior judgement" becomes "uses AI well." Everyone signals the same way. The signal collapses.
4. **Recruiters and capital lose the ability to read labour.** CVs blend together. LinkedIn posts converge. Interviews are suspect.
5. **Capital, faced with an unreadable labour market, leans harder on infrastructure deals.** It buys AI deployment vehicles instead of betting on individual hires. Back to step 1.

Each step is rational on its own. The aggregate is a market that **gets faster at moving money and slower at moving people.**

This is not a labour-market reshuffle. It is a **signalling-system collapse**. And it explains the one fact that the doom-vs-optimism debate cannot accommodate: AI agent benchmarks are failing in 2026, but AI is reshaping the job market anyway. Capability is not the variable. Signal legibility is.

Consequences

Four that impact executives.

Firms hollow out their own judgement bench while believing they are upgrading it. They deploy AI at the capital layer (Story 1). They drop juniors (Story 2). They screen for AI fluency on the way in (Story 3). The result is a workforce of polished AI-users who lack the judgement that AI itself cannot supply. The competence gap is invisible until a real edge case arrives, a regulatory inspection, a client crisis, a system failure, and the bench has no one on it who has actually navigated one before.

Capital allocation routes around labour markets that no longer perform price discovery. PE-financed AI deployment proliferates because it bypasses a hiring market that has stopped giving readable signals. This is rational from the capital side and corrosive from the labour side. Wages, mobility and matching all degrade as a downstream effect.

The mid-sized non-portfolio firm becomes the structural loser. It cannot read its own applicants, cannot afford to lose its juniors prematurely, and is not on the receiving end of PE-financed AI deployment. It carries the costs of all three stories without access to the capital that makes them survivable. If you run such a firm, I believe that this is your single most important strategic problem this year.

Trust becomes the only remaining currency. When CVs cannot be trusted, when interviews cannot be trusted, when AI deployments cannot be trusted to perform at the rate the deal terms imply, what is left is referral, reputation and demonstrated track record outside the AI-mediated channel. This is what Gilman's "judgement economy" looks like in practice, not a soft humanism, but a hard market response to signal collapse.

Enrique Dans put the corporate corollary best in his March 2026 Medium piece: "*AI does not create strategic clarity. It reveals whether you had any to begin with.*" The same holds for hiring. AI does not produce a better workforce. It exposes which firms had the judgement to recognise good people in the first place.

What this means for your work

I have spent a career on the corporate side of these decisions, first as a PMO director and now building AI products. The pattern I have seen play out in firm after firm is the same. **The advantage is not having AI; every firm will have AI inside twenty-four months. The advantage is the discipline of building the constraints and signals that AI cannot supply on its own.**

Three actions, one each for the audiences who will read this.

For executives. Stop treating AI deployment as an IT decision. Treat it as a **signalling-system decision**. Audit your hiring process for what signals you can still trust. The Architectural Decision Record discipline I have written about elsewhere, used to make AI deployments auditable, applies just as cleanly to hiring. Write down the decisions that survive AI mediation; protect them; review them. If your hiring process has no readable signal that survives AI on both sides, you do not have a hiring process. You have a queue.

For mid-career professionals. The instinct to signal AI fluency on LinkedIn is correct in the short run and corrosive in the long run. Everyone is doing it; the signal is dying. Invest instead in the things AI cannot mediate: live judgement under uncertainty, demonstrated track record on a project a senior person will vouch for, the kind of cross-functional fluency that a BA or PM acquires by holding contradictory stakeholder demands in their head. Those are the signals that survive when CV polish becomes worthless.

For graduates. Apply less, network more. The application channel is broken. Greenhouse can tell you exactly how broken. The channels that still work, referral, in-person events, direct outreach to specific operators in specific roles, are the ones AI cannot saturate. They are slower. They are also the only ones whose signal-to-noise ratio is still legible. Choose your battles accordingly, and do not internalise the rejection rate of a system that is no longer functioning as a market.

Closing

The familiar AI debate poses a binary: are the machines coming for us, or are they freeing us? Both framings miss the mechanism. The mechanism is signalling. Signals fail before jobs do. Capital reads faster than

labour because capital does not depend on the same signals to function, it can route around them, buy infrastructure, deploy through portfolio companies. Labour is stuck with the broken market.

Until the signalling system is rebuilt, by firms, by professionals, by the institutions that train graduates, the loop continues. Capital deploys, work redistributes upward, workers signal AI fluency, signals collapse, capital deploys harder.

The story to watch in the next twelve months is not whether AI replaces white-collar workers. It is **which firms rebuild a hiring process they can still trust**, and which ones simply outsource the decision to capital and find, five years from now, that they no longer have a senior bench worth deploying AI on top of.

Before AI takes anyone's job, it takes the market's ability to tell who is good at the job. That is the part nobody is pricing yet.

Key takeaways

- **Capital is setting the speed of AI adoption, not productivity.** PE-financed deployment vehicles like the Anthropic–Blackstone–Goldman venture are pushing AI into mid-market firms ahead of any capability proof.
- **AI is failing white-collar benchmarks and reshaping white-collar work simultaneously.** The disruption is vertical: the bottom rung disappears; the top rung becomes more demanding; the apprenticeship pipeline collapses.
- **Both sides of the labour market are now AI-mediated.** Candidates apply through AI; recruiters screen with AI; the resulting signal collapse pushes capital to route around the labour market entirely.
- **The mid-sized non-portfolio firm is the structural loser.** It carries the costs of all three dynamics without access to the capital that makes them survivable.
- **The variable to manage is signal legibility, not AI capability.** Firms that rebuild trustable hiring signals will retain a senior bench. Firms that do not will run on borrowed judgement until the bench runs out.

Sources cited

- *Financial Times*, "Anthropic teams up with Blackstone and Goldman Sachs in AI consulting venture," news segment with Ryan McMorrow, 5 May 2026.
- *Perplexity*, "Anthropic partners with xAI for computing power to ease Claude capacity crunch," 7 May 2026, <https://www.perplexity.ai/page/anthropic-partners-with-xai-fo-Hnp7fTXLSSyCrA6BBozT.g>
- Nils Gilman, *Noema*, "How To Future-Proof Your Career In The Age Of AI," 9 April 2026, <https://www.noemamag.com/why-a-liberal-arts-education-will-soon-be-more-valuable-than-ever/>
- *The Guardian*, "The big AI job swap, why white-collar workers are ditching their careers," 11 February 2026, <https://www.theguardian.com/technology/2026/feb/11/big-ai-job-swap-white-collar-workers-ditching-their-careers>
- Joe Procopio, *Entrepreneurship Handbook*, "It Turns Out, AI Agents Suck At Replacing White-Collar Workers," 6 February 2026, <https://ehandbook.com/it-turns-out-ai-agents-suck-at-replacing-white-collar-workers-8f2657ca61a0>
- *TechCrunch*, "Are AI agents ready for the workplace? A new benchmark raises doubts," 22 January 2026, <https://techcrunch.com/2026/01/22/are-ai-agents-ready-for-the-workplace-a-new-benchmark-raises-doubts/>
- *Business Insider*, "Most Anthropic Teams Are Coding With AI, but It's Not Replacing Engineers," October 2025, <https://www.businessinsider.com/most-anthropic-teams-coding-with-claude-ai-not-replacing-humans-2025-10>

- King's College London, "New study reveals early impact of AI on job market in UK," October 2025, <https://www.kcl.ac.uk/news/new-study-reveals-early-impact-of-ai-on-job-market-in-uk>
- Efosa Udinmwun, *TechRadar*, "From finance to tech, AI is quietly rewriting LinkedIn posts," 18 February 2026, <https://www.techradar.com/pro/its-not-just-you-linkedin-is-being-flooded-by-awful-ai-generated-posts-and-tech-workers-are-particularly-bad-at-it>
- Publicate, "The industries using AI the most on LinkedIn," 2026, <https://publicate.it/blog/the-industries-using-ai-the-most-on-linkedin>
- *Financial Times*, FT analysis on graduate unemployment with management editor Angeli Raval and Greenhouse CEO Daniel Chet, 20 February 2026.
- Enrique Dans, *Medium / Fast Company*, "AI won't replace strategy: it will expose it," 15 March 2026, <https://medium.com/enrique-dans/ai-wont-replace-strategy-it-will-expose-it-05035dbfbce3>
- *MarketWatch*, "Software ate the world. Now, Wall Street is worried AI will eat software," 4 February 2026.
- *Perplexity*, "Anthropic AI tool launch triggers SaaSocalypse; Tech CEOs urge calm," 9 February 2026, <https://www.perplexity.ai/page/anthropic-ai-tool-launch-trigg-L42EqhJZTFqSJ9x9PpCAUQ>
- Wences Alfageme, "Stop Hiring AI Engineers. Train Your BAs and PMs Instead," 27 April 2026, alfafintec.com/portfolio.html

[#AIStrategy](#) [#FutureOfWork](#) [#JobMarket](#) [#HumanCapital](#) [#AIGovernance](#) [#Hiring](#) [#JudgementEconomy](#)
[#ProjectManagement](#) [#BusinessAnalysis](#) [#DigitalTransformation](#)

Disclaimer and Disclosure

Third-party Content and AI Assistance: This article references publications, studies and tools owned by their respective publishers and creators. The author does not claim ownership or affiliation with these third parties. This article was written by the author with assistance from Generative AI Language Models.

Transparency Notice: While every effort has been made to ensure accuracy, readers should verify information independently and consult original sources where decisions of consequence are involved. The use of AI in the writing process is disclosed in the interest of transparency, but all opinions and analyses are the author's own unless otherwise stated.