

The Great AI Talent Migration: Why Universities No Longer Anchor Frontier AI Research

The Economy Research Editorial^{1,2}

¹The Economy Research, 71 Lower Baggot Street, Dublin 2, Co. Dublin, D02 P593, Ireland

²Swiss Institute of Artificial Intelligence, Chaltenbodenstrasse 26, 8834 Schindellegi, Schwyz, Switzerland

Abstract

This article considers the migration of AI talent from the university to industry as a structural change in knowledge, research and the creation of human capital. We argue that the development of large-scale AI, particularly generative AI, has neither simply brought a significant new instrument to the classroom and the lab nor that it has changed which organizations are at the research frontier nor who controls the frontier. The expansion of capital intensive model development, access to proprietary datasets and astronomically high wages now place firms, not universities, at the heart of cutting-edge AI research. Yet this movement into industry involves much more than shifts in staffing of university labs. The article considers how the migration also changes student incentives, shrinks the potential for the production of free scientific knowledge, exacerbates inequality, and alters education-labor market links. We make the case that for all its normative appeal, full recapturing of the research frontier by the university under current conditions is materially unfeasible; a plausible alternative path involves leveraging university strengths in fundamental training, transparent signaling, openness of inquiry, ethical reflection and social study of AI's consequences. The article suggests that policy should move beyond attempts to reinstate a lost academic order and redefine the university's place in a knowledge system increasingly structured by industry.

1 Introduction - The Structural Exit of AI Talent from Academia

The advent of AI is often celebrated as a paradigm shift that will turbocharge education and productivity, extend human capacities and relieve the burden of research. But underneath that rosy picture is a much more fundamental change, which is the departure of talent and innovation from the ivory towers and into the corporate halls—causing a tectonic shift to the architecture of knowledge generation and skill development. To frame AI simply as a new tool akin to a calculator, and the internet is grossly inaccurate: that framing misses the stark physical differences of the real thing. State-of-the-art AI—particularly large generative models—boasts computing, data, and expertise requirements that most are not prepared to provide, requiring multi-million-dollar teams to build.^[1] At present, the innovation ecosystem is profoundly lopsided: according to new research by Stanford HAI, no university yet has the compute to match Google’s state-of-the-art Gemini Ultra, whose compute cost exceeds the annual spend of the top research university.^[2]

The common wisdom about AI in education treats it as a kind of basic tool that makes the distribution of knowledge easier, without radically changing the players involved in the game. This analysis shows that the thesis is also incomplete. As data showing the flow of AI specialists from universities into major tech firms rises^[3] and as the salaries of AI scientists soar to six and seven figures,^[4] a systemic dynamic is playing out that will shape the nature of human capital and inequality over the coming decades. It will shift taxpayer-funded resources and incentive structures to private businesses and away from public teaching and research institutions. It will change the nature of the human capital in education, how it is learned and to what extent it brings economic benefit; as a result, it is worth it. Essentially, AI is no longer just an assistant, but an active agent that shapes learning and employment. For example, in 2019, the share of U.S./Canadian-based AI scientists in private industry surpassed 68%; this marks a five-point increase from just a decade or so prior, when industry employed roughly 48% of AI talent.^[5] The over-enrollment in computer science at elite universities and the rolling in of an unsettlingly large number of high–six to seven-figure salaries has completely upended the classic academic pathways and dashed expectations of research-heavy classes. When the generative AI boom hit in 2023, however, it only magnified this trend, with reports showing that the proportion of U.S. and Canadian AI PhDs that joined industry skyrocketed by a surprising five and a half percentage points from 2021 to 2022, reaching almost 71%,^[6] while more and more academic tech treats declined to apply for tenure tracks. Meanwhile, academia balances how to teach people about AI with how to let research carry on as usual.

This article proceeds in four steps. Section 2 traces the early movement of AI research toward industry. Section 3 shows how this shift intensified in the era of deep learning, large language models, and generative AI. Section 4 examines the consequences for public knowledge production and the declining centrality of universities. Section 5 considers whether universities should attempt to reclaim the center of AI research or redefine their role more selectively. The conclusion then draws out the policy implications.

The article examines the significant time required for second-order effects on learners or society at large, as opposed to simple input-output efficiencies or gains. How does having access to an infinite supply of AI tools shape student learning and teacher evaluation? Will weighing AI tools more heavily in the learning process undermine critical thinking? How will the concentration of AI research labs in corporate firms affect knowledge

dissemination and inequality? We motivate these questions with existing evidence: Stanford's study utilizing AI in education indicates that AI only improves performance temporarily, yet those improvements fall away as soon as the tools are removed.^[7] Brookings has expressed concerns around 'cognitive stunting' in scenarios where students may overuse AI to think. Evidence from the World Bank^[8] and OECD^[9] points to AI's ability to replace routine roles and favor high-skilled AI talent. Post-ChatGPT, jobs requiring high AI susceptibility fell 12,18%,^[10] especially in (often entry level or administrative) roles with greater AI susceptibility; such impacts will likely exacerbate wage inequality: in the OECD study, 0.3% of the workforce combined accounted for AI-skilled individuals, who disproportionately occupied top deciles, earning over 10% more than the next best offer.^[11]

Using this evidence with a causal logic, the evidence suggests how the migration of AI talent systematically benefits human capital development, institutional design and inequality through policy implications. We then translate our causal analysis into a series of policy recommendations. In section 6, we discuss curriculum and pedagogy implications (AI literature, critical thinking skills, deep learning, focused tests of knowledge), institutional implications for administrators (ethics guidelines, training programs for staff, institutional planning) and labor market and education implications for policymakers (funding campus AI research, reforming education labor market regulations). We examine counterarguments, like those claiming that AI democratizes knowledge or just keeps up with technological trends, using evidence-based reasoning to reinforce our main claim. The conclusion summarizes our argument and makes a principled call to action.

2 Early Industry Transition

The migration of AI talent out of the academy and into the corporate world truly became endemic in the 2010s, indicating a decisive change of course that is a mirror of current trends. While the breakthroughs in the early years of AI came from academic institutions such as Stanford, MIT and the University of Toronto, by 2012–2014, tech heavy hitters like Google (Google Brain and DeepMind), Facebook (FAIR) and Microsoft (Research) entered the game in earnest, rolling out major AI research programs. Their data, hardware, money and salaries accelerated the competition for academic talent. A simple survey, for instance, suggests that by the early 2010s, the overwhelming majority of AI talent was working for private players, peaking around 2019 at 68% of AI researchers in industry from just 48% in the early 2000s.^[12] The Stanford HAI AI Index shows a similar trend: 70.7% of new AI PhD graduates have selected an industrial position, up from around 41% for the United States and Canada in 2011.^[13] Private rivals won, once again.

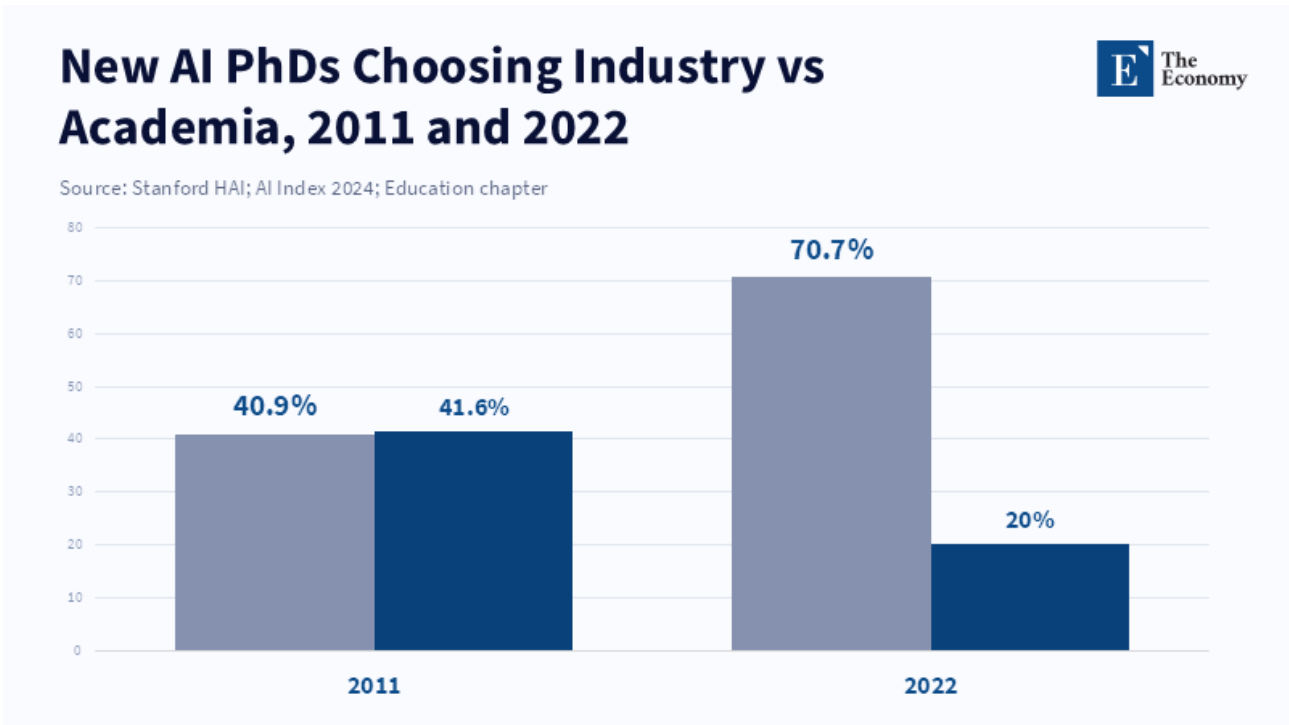


Figure 1: The doctoral pipeline no longer replenishes universities; it feeds firms

There are many intertwining causes for this trend, but the most compelling was probably simply monetary. Silicon Valley and its equivalent were paying AI specialists far higher than university salaries—Akcigit et al. put it quite succinctly: 'top AI salaries in industry increased from \$595K to \$1.9M' [14] (meanwhile, students still earn less than the same). Going into private industry, of course, pays an offset too, but even excluding the superstar, Salary, idols, PhD students and postdocs earned way higher lifetime earnings in industrial work than on the tenure track. Said one contributing member of the industry: 'The huge salary disparity between industry and academia was a major motivator.'




Figure 2: The pay gap is not marginal; it is large enough to redirect careers.

Beyond pay, it was also about access to resources and career opportunities. The ever-growing computational power and data resources that firms could tap into made projects feasible that would be impossible for academic laboratories. For example, the pioneering of deep neural networks required months of GPU training for prolific internet data and as soon as technology giants proved to the world that they could do so and produce great results, talented young scientists knew they had to be working in such a place if they wanted to make the greatest impact. The RFBerlin discussion paper finds that getting access to major AI conferences was a big step in getting recruited. Being accepted to a publication at a leading AI conference increased the odds of obtaining a position at a top firm by 2 to 6%. In other words, students for whom groundbreaking research was published in the blink of an eye found that they could suddenly start working in industry. Impact and prestige were other motivations; industry labs were fast becoming the motors of AI.

A global context that sped up the initial transition should be included as well. World-class talent was starting to move across borders, with North American, European and Asian tech hubs fighting over a share. For instance, the 2025 Global Talent Tracker report indicates that despite the US still being the top destination, the UAE and Saudi Arabia's year-by-year position has improved markedly. Regarding supply, Indian technical universities have been training many emerging country AI engineers, most of whom are immigrants and one Indian technical university has been identified as the largest single source for industry, destined for internationally mobile AI talent. Hence, universities from developing countries started populating the increasingly industry-centric world ecosystem.

This long-term result was slow, but significant. As the pipeline of opportunity in industry grew and grew, universities began to respond, expanding AI-related courses and research centers to try to prevent the talent from flowing to industry, but they could never match the incentives industry offered. Graduate students would invest more in skills, building that would look good on their resumes to industry, faculty hiring committees would value applied publications more highly and AI-related programs in particular grew in popularity, largely due to the money to be made in tech jobs.^[15] Eventually, universities' centrality to AI innovation would fade and student decisions would follow: over the long run, this process will alter the training and supply of AI talent. If talented students perceive that they will work in industry, they will go with the modalities that lead to industry at higher rates than the more theoretical ones.

Finally, that first change created the foundation for what has since become a reshaping of the AI ecosystem. Over the course of the 2000s and 2010s, growing investment and human resources created a mostly, private AI system, made even more so due to tech firms' advantages in time and money and access to data and supercomputers, relative to academia (evidenced by the lack of faculty at AI conferences and the amount of work that is industry, related, industry, originated, or is co, authored by academic researchers.) This stems from a scholar, industry hybrid that created a new business, driven research cartel and made academia more sensitive to industry demands. The following section highlights how pervasive megatrends—like generative AI and zero-sum competition for talent—have turned a stagnant process into a rapid one with mass implications.

3 The Trend Got Much Stronger Recently

The migration of AI talent has accelerated dramatically over the last 3 years, spurred by generative AI innovations as well as increased competition. The release of LLMs such as ChatGPT in late 2022 was yet another inflection point, as new applications emerged overnight and companies rushed to hire more AI talent. These figures speak to the increased pace: in Stanford’s 2024 AI Index, the share of new PhDs in AI joining industry jumped by 5pp between 2021 and 2022^[16] (to 71% in North America) in tandem with academic opportunities shrinking by a similar proportion. This 1-year escalation was unprecedented in the recent history of academia, where the share of PhDs in industry rose steadily in the prior decade.




Figure 3: The outflow to industry is rising just as the return flow into academia weakens.

These sharp deviations are explained by three elements: a mountain of investment in both compute and capital by the biggest names in the technology industry, a fundamental evolution in the nature of research in generative AI and a change in institutional policy, the study environment, and students’ perceptions of AI. First, competition for both talent and compute intensified massively. Gemini training by Google alone will have used \$100, \$200 million in computing in 2023,^[17] Meta dangled nearly \$100 million in signing bonuses^[18] in front of potential AI researchers and investors mushroomed as wealthy funders sought their piece of AI prosperity. These signals provided faculty and students long-term incentives that big technology companies would pay anything to win the scramble for talent and dominance—offering financial rewards and the prospect of being launched on the frontier of AI research.

Second, the nature of the research environment changed significantly. Generative AI has provided dramatic leaps in capabilities—sometimes approaching or surpassing human performance in language, vision and tools—whereas prior deep learning progress was more steady. These breakthroughs attracted heavy media coverage and investments, attracting not just engineers but VC investors, aspiring researchers and curious ama-

teurs, leading to a veritable explosion of start-ups and elite corporate research. One industry observer identified 51 major ML models announced by industry labs in 2023, but only 15 from academia^[19] (who can't keep up with the scale of industry models anymore [50 times bigger, per Stanford]). In the old system, academics built the big models, but now they are working off of pre-trained, closed, prebuilt models, representing a qualitative shift in the innovation process from open to closed. This new influx of researchers into industry creates a virtuous cycle—the more successful industrial AI is, the more talented academics it draws, which then spurs more powerful systems to be built.

Third, government and institutional drivers are impacting research talent pathways. Many nations view AI as a high-priority area and are opening their gates, funding top research centers and streamlining their visa processes. Student and researcher visas for AI have been fast-tracked in the US and Canada; Europe and Asia have announced ambitious national AI initiatives. But even academics are experiencing conflicting pressures: while hiring in industry has surged, new tenure-track positions have not and students have less incentive to view academia as a finishing school. One comprehensive Lumina, Gallup 2025 poll noted that, although degrees are appreciated, graduating students have described their curriculum as being quite disconnected from industrial demand and that 45% of students claim their classes do not have defined rules for AI use,^[20] implying that AI is no longer out on the fringes but in the classroom. Google's set of generative AI has resulted in students having familiarity with AI, thus confirming that it is a vital tool and a foundational life skill.

These recent trends have done nothing but continue and intensify the pre-existing effects: the advancement of generative AI has seen to it that industry roles are seen to constitute not only greater financial benefits, but also potentially greater capacity to impact the field, as an academic commented: "by 2025, the environment for AI research in academia was effectively a vendor,customer relationship with frontier labs," where academia mostly works the industry blueprints on account of an exorbitant expense required to build large systems from scratch. This feedback loop attracts more and more researchers to a handful of tech major hubs where their work will be used to serve closed and proprietary innovation cycles.

One manifestation of this cycle can be observed via research publications. When a researcher moves to industry, publication output drops drastically (studies quote approximately 1000 fewer AI publications and 2000 fewer CS publications annually because of the drain of talent from research to industry) and the number of patent applications rises close to 6x,^[21] demonstrating a real change from open academic sharing to closed IP production. It is this "moat" of research installed by talent and infrastructure that raises switching costs and slows the dissemination of new insights in a privately held environment.




Figure 4: Talent migration reduces open science, not only campus headcount

A rising trend can also impact the labor markets generally. Analysis of US job postings by the World Bank shows that, after the public debut of generative AI, demand for jobs that would be susceptible to automation fell by 12% and in mid,2025, had further fallen by 18%^[22] relative to other jobs. Entry jobs (18%) and unspecialized jobs (20%) were disproportionately harmed; specializations such as administrative assistance (40%) and professional services (30%) were particularly harmed, with an immediate impact on employment. But demand for highly skilled workers was not showing signs of surging demand, while salaries commanded a premium for their labor. Such a reversion is a sharp departure from previous waves of automation, where adoption was a gradual process. Here, adoption has been erratic and, with document automation tools across finance and law firms, as well as AI assistants across companies, is displacing the work of less experienced analysts and assistants.

Universities feel the pressure of this new environment too, but not necessarily with respect to talent migration—massive generative AI adoption has also ushered in the end of minds’ adherence to rules. As an example of this latter effect, academia’s larger culture has begun to summon a high-ceiling discussion of academic integrity in the face of the widespread use of generative AI. Lumina, Gallup found that more than half of colleges forbid AI use in class work,^[23] yet students’ practice repeats itself—the rules forbid the use of AI, but students use it. Universities’ dilemma, then, is whether to not use AI or to use it, but teach students how; early research on this subject points towards possible degradation of learning: AI writing tools performed better on a writing task than students with no tool, but their improvements disappeared when the tool was removed,^[24] allowing students to write better answers but not learn more.

In sum, the recent period is just this pre-existing dynamic writ large. The Boston dynamism of the 2010s was amplified in 2023–2025 by the acceleration that generative AI brought with it. As the industry is now taking on even more voracious hiring in comparison to the university; as governments are now engaging competitively in a

race to the top in AI funding; and as the speed of technological evolution means one can barely be at the frontier if one does not have an industry, scale research team, then effectively teachers and students (navigating the challenge of tool proliferation and curriculum disruption) and professors (watching as the academic infrastructure of research they had been overseeing collapses under their eyes) have all been caught up in something new. It should be obvious that what I have described is a series of long-term problems. What follows is the consequence: the university’s age in the leadership of AI research is long gone.

4 University No Longer Is the Center of Research

The fast pace of private sector progress and the presence of industry in all parts of the tech life cycle, has pushed universities away from a dominant position in front-edge AI that they occupied in the past. This is not solely quantitative—industry centers far outpace university centers in terms of progress; according to the Stanford HAI policy brief on public sector AI, a lone university ran only 15 machine learning projects in 2023 relative to 51 by industry,^[25] while government research only undertook two.

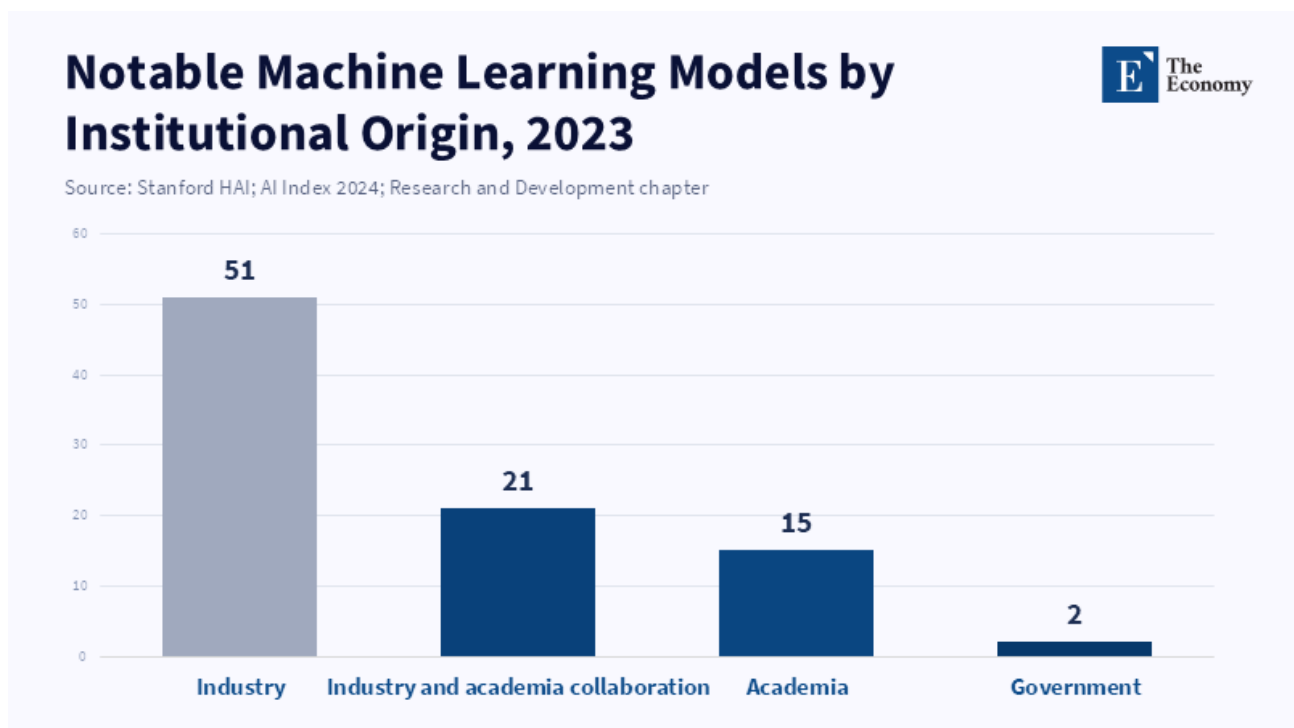


Figure 5: Frontier model production is now overwhelmingly corporate.

Compared to a decade ago, when research produced a roughly equal number of important ML models, now the private sector has a substantial majority of models. This gap has grown even wider since, in 2023, Stanford reports the private sector had spent \$300 billion dollars on AI research since 2014,^[26] while university spending has either now plateaued or declined in real dollars in recent years. Private sector researchers are also reported to now produce more than 1,000 times the amount of capacity compared to all university centers combined. Due to this disparity in scale, all but the most elite universities are unable to receive ‘run the world’ projects and according to one authority, the scale of hyperscale computation necessary for state-of-the-art research is unfeasible not only for the top university but for a university alliance. Academic labs lack the capacity to run

the complex models that require an immense amount of data and processing power.

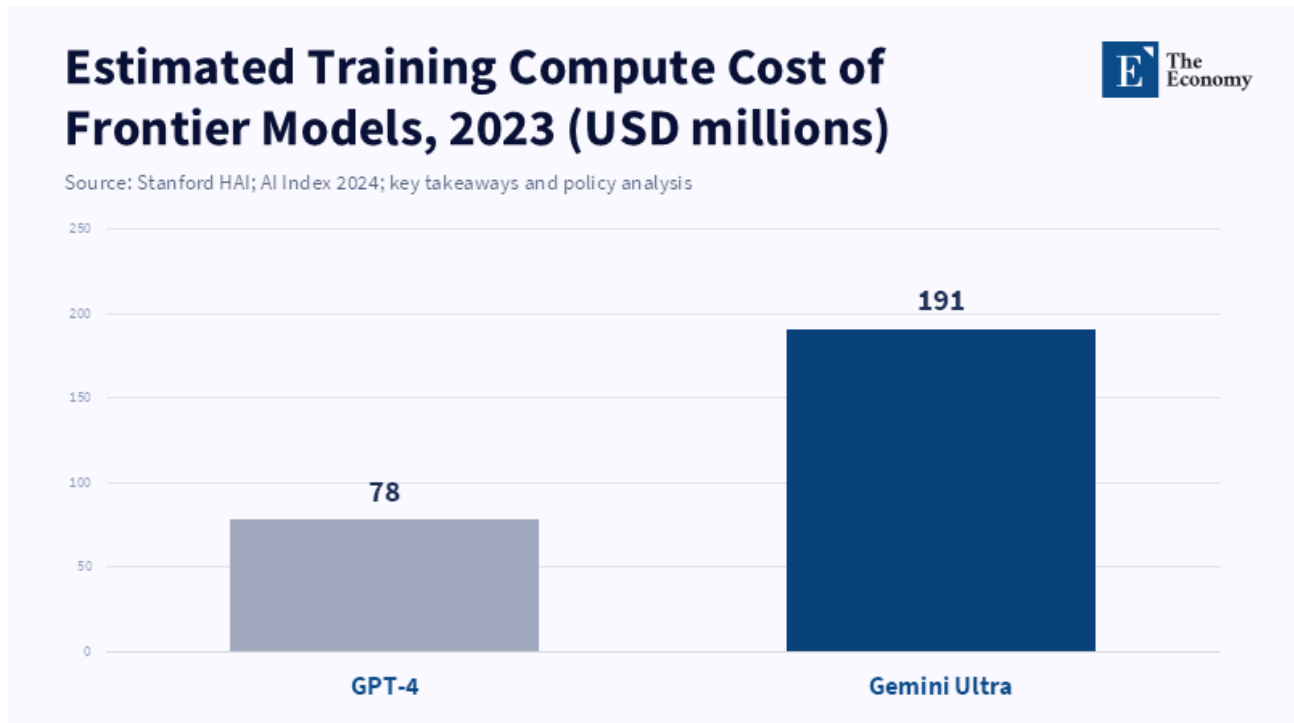


Figure 6: Once frontier training costs reach eight or nine figures, universities stop being realistic frontier builders.

This change in funding has completely changed incentives for the universities. As the industry is now in charge of the training of foundational models, more and more academics have shifted towards a narrower research while favoring theory, algorithm efficiency and certain applications and very rarely doing research that involves large models with hundreds of millions of parameters. This is reflected in the number of publicly available open-sourced papers from university researchers. When considering those who switch to big tech companies, the number of public papers for a researcher drops drastically, according to an RFBerlin report.

Summing up the changes, estimates that the net effect of these positions is that 1000 fewer papers are being published a year now. Again, this figure is balanced out by the increase in patent submissions, showing that knowledge is being hidden and still getting to the industry in an indirect way. A researcher will maintain 6 times the number of patents^[27] in a company than they will in a university on average. This revelation from open science to proprietary research and development has huge implications. For instance, university science is made wide and available while industry is largely hidden away, hence a significant amount of other knowledge remains hidden behind "pay walls", thus there is little way we can make much progress in the fundamental science of AI.

This erosion of power in academia has had a tremendous impact on students and curricula within universities. Courses in AI and ML increasingly tend to teach the application and use of frameworks and tools rather than the theoretical background and the underlying principles of their operation.^[28] Instead of being taught to code up an AI model from scratch, professors tend to rely on industry-oriented practical examples or have guest lecturers from big tech companies teach classes. More often than not, in classes on these topics, the discussions are led by students who are working with one of a handful of major tech companies. Given this focus, research skills are built around proprietary data sources and tools, which can be damaging, though there are some

worthwhile advantages to it (food, drugs, data, etc.). These advanced courses are increasingly used by students as a stepping stone into industry internships and more and more students are taking these courses not as an academic experience, but because they are expecting to gain entry into a high-paying industrial career. This has created very strong incentive systems in the universities, which value corporate collaborations and papers published in high-impact industry venues over academic research. Universities nowadays aggressively seek these partnerships, substantial grants and chairs in academia and so forth, which might ultimately detract from the university's academic goals. This shift has also changed prestige structures in the area, with the "edge" being defined by the presence of a partner industry and a gigantic hardware budget behind your work. While universities can contend with this by establishing multiple university research centers and lobbying for public funding for AI research workshops, the projects will never reach the same scale and level of development as those of big tech companies.

Perhaps the most significant consequence of this huge change in research ownership and productivity is the way knowledge is generated, circulated and utilized. For previous decades the broad and open reach of universities in AI research and the one-of-a-kind opportunity to share new ideas and concepts across the world was unparalleled. But with this recent surge of intense skill concentration within a handful of domestic tech businesses, the arrival of new ideas might be more sluggish compared to what it previously was, if any. In fact, this ultimately causes more inequality; the knowledge generated inside these private companies is funneled into a competitive advantage and to profits and the World Bank and OECD have transparently provided evidence that the workers in the AI field are already at the top end of the pay scale and those in receipt of a pay premium for AI skills generally already occupy the highest, paid jobs in the country, which in turn, with this sharp surge of concentrated skill inside a handful of private businesses, contributes directly to the widening income gap, both within nations among the select few in the private sector and between the dominating businesses in high, income countries and other less, redeemed organizations in both academia and emerging nations.^[29]

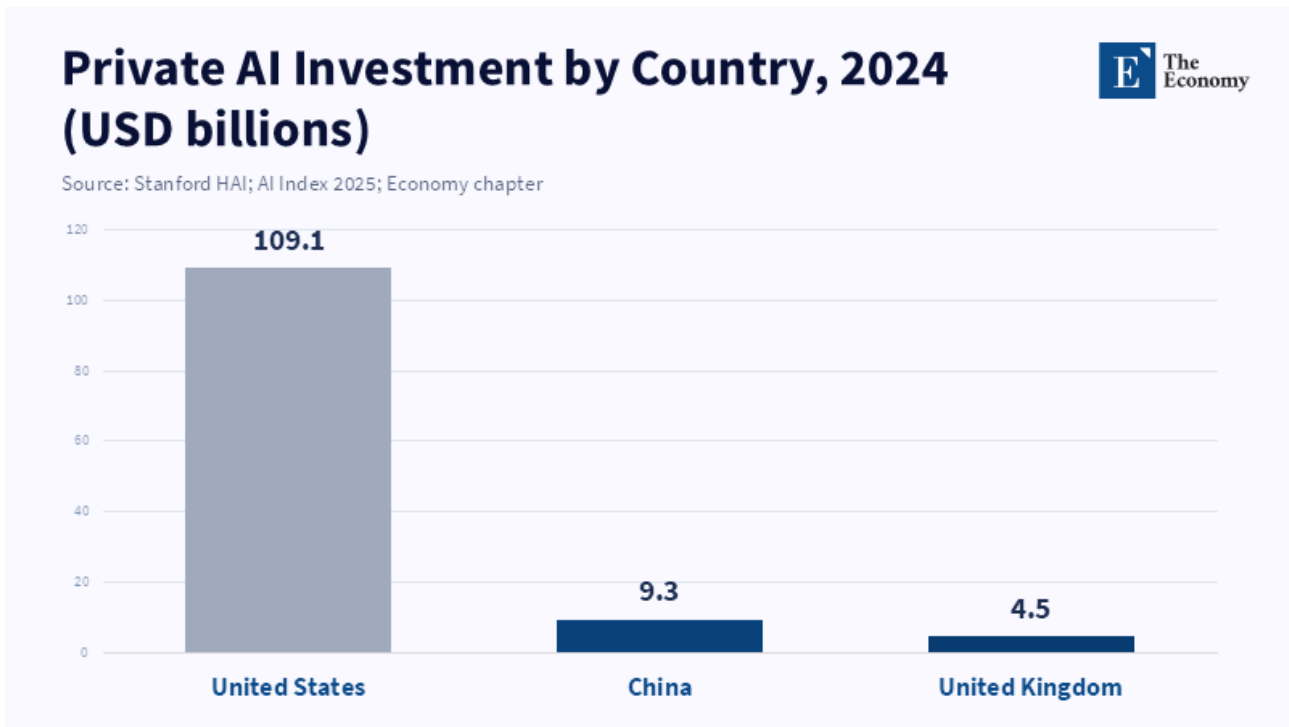


Figure 7: AI capital is geographically concentrated, which reinforces research concentration.

All evidence suggests that the once central role of academia in AI research has been diminished. The rise in corporate control over the field translates to ramifications for knowledge production and popularization. Universities have moved from innovators and codifiers of advanced technology development and distribution, to consumers of corporate technologies and training plants for "private sector" employees. These turn of events mean that universities are being called to assume a new role; that of advising ethically, reflecting on the impact of technologies and training people with the "basic job skills" not requiring advanced AI knowledge. And most importantly, these circumstances force us to radically overhaul the expectations and role of the university. Should they then follow the much more computationally intensive regimes of research, or should they remake their actual role? That is the issue of the next paragraph.

5 Why or Why Not Do We Need Universities to Come Back to the Center of Research?

Given the stated leadership of the sector, the tabog answer to this type of policy question is "yes". As suggested by the plans of the Stanford Institute for Human-Centered AI, academia, as the host of transparency, in extremis and non-profit motivations, can draw on historically advantageous starting positions to order the pathway for AI's societal implications. They note that academia has historically invented the defining innovations for AI, from backpropagation to open source tools, "and it should lead the way again in guiding responsible innovation". Their response called for increased public investments in universities, based on AI research, the development of open supercomputing infrastructure and the promotion of dedicated research and policy efforts around a unique type of AI: safety, ethics and education. They reason that without academia, industry remains the only AI engine and will direct AI development toward only those applications deemed "sufficiently profitable", while

neglecting vital considerations like fairness, interpretability and broader societal advances in domains such as medicine and climate sciences. Such "coming back" would manifest as a network of campus-led public sector AI research programs, funded by philanthropic and governmental contributions.

There are many strong reasons against such a proposal, however. The first and most obvious in terms of practicality: as discussed in the above section, the scale of infrastructure needed for state-of-the-art AI is mammoth. If today's best industry R&D labs are willing to pour hundreds of millions of dollars (rather than pounds or euros) into one project,^[30] how would publicly funded universities, constrained by budgets and university governance, possibly compete with this? Even assuming large emergency funding as mentioned in the previous section, the pace of bureaucracy on the public sector side vastly exceeds that found in industry. A researcher used this anecdotal quote in a similar context: "There is no frontier AI research that can be carried out at universities. . . [they] are blocked off by the massive resources needed for hyperscale computing". Anecdotal evidence from PhD students seems to support this. Apart from long bureaucratic wait times for computing cluster resources, the offer from industry partners of priority access could just push talent towards the private sector laboratories.

Given these effects, is there any reason at all to hold on to the old paradigm of universities leading the delivery of artificial intelligence? Based on the paper's findings, it is unlikely; we would argue for a more measured approach. If policy discussions persist with attempting a direct top-down replication of the industrial model, then this is liable to lead university resources toward developing large language models and away from those areas where universities can add the most value. Universities can better focus on those areas in which they have the strongest comparative advantage: first, an emphasis on fundamental training, such as teaching students the mathematical, statistical and conceptual foundations of AI;^[31] 2 second, a focus on longer-term, blue-sky issues, such as trying to determine the boundaries of the capabilities of learning algorithms, or to conducting interdisciplinary investigations into the boundaries of AI and the social sciences and humanities;³ and third, a focus on areas that won't necessarily be high priorities for consumer and industry adoption:⁴ the advancement of basic research, for example into the theory of learning algorithms or into neuro-symbolic AI and the conduct of discussions around the broad ethical and societal impacts of this technology.

Universities also have a second unique function, bearing certification and signaling competency. Degrees symbolize a minimum threshold of credibility. They can emerge as the market's certification agent and mass medium for such competency, not by competing with industry in hardware, but by focusing on providing clear, credible paths to certify talent. This would mean demanding extremely high standards in AI-related degrees, so that jobs requiring university-level, related time and effort will have candidates with high intellectual capacity. It could mean deemphasizing certain curricula, demanding higher levels of skills, including those that remain nontoothless to machine learning—the ability to think critically, creatively and morally with a high level of capacity. These are less likely to be automated and could be highly relevant for many jobs. If universities indeed become the natural "accrediting institution" for AI talent, they will continue to have a voice in the ecosystem even if they don't push the limits of AI engineering.

Fields outside of the core disciplines will also need to be addressed, such as the schools of education and cognitive science that can probe more deeply into the implications of AI for education and cognition^[32]—things

that a product-focused industrial research lab would have little interest in pursuing. These departments can also pave the way in creating and evaluating innovative pedagogy that can be subsequently scaled to help inform a generalized method of teaching (e.g., following the learning pathways of students over a decade with and without personalization aided by AI). Other fields, such as those found in public policy and ethics, can direct research towards AI governance, rights to the data and socio-technical consequences. The economies of scale inherent in studying a problem over decades may not appeal to corporate labs, but may appeal to a university setting. Academic institutions should be rewarded by governments to do this by supporting inter-institutional research on important cases.

In summary, there is no reason that some of the "AI as a public good" discussion needs to take place entirely within the industry. The market alone may overlook broad societal consequences and academics should recognize their role in AI "fundamentals" gathering, open science and partnerships with industry focused on building the foundation for industrial-scale research. The university's role may need additional policy tools, such as federal funds for "public" AI platforms, providing academics with throughput and storage (analogous to the Stanford proposal for shared supercomputing time) and regulation that results in public access to data sets or research. At the same time, academic institutions will need to enforce stronger standards of academic integrity, ensuring AI tools in the classroom are used to learn and not just for gaming the system and providing policy guidance on their use.^[33]

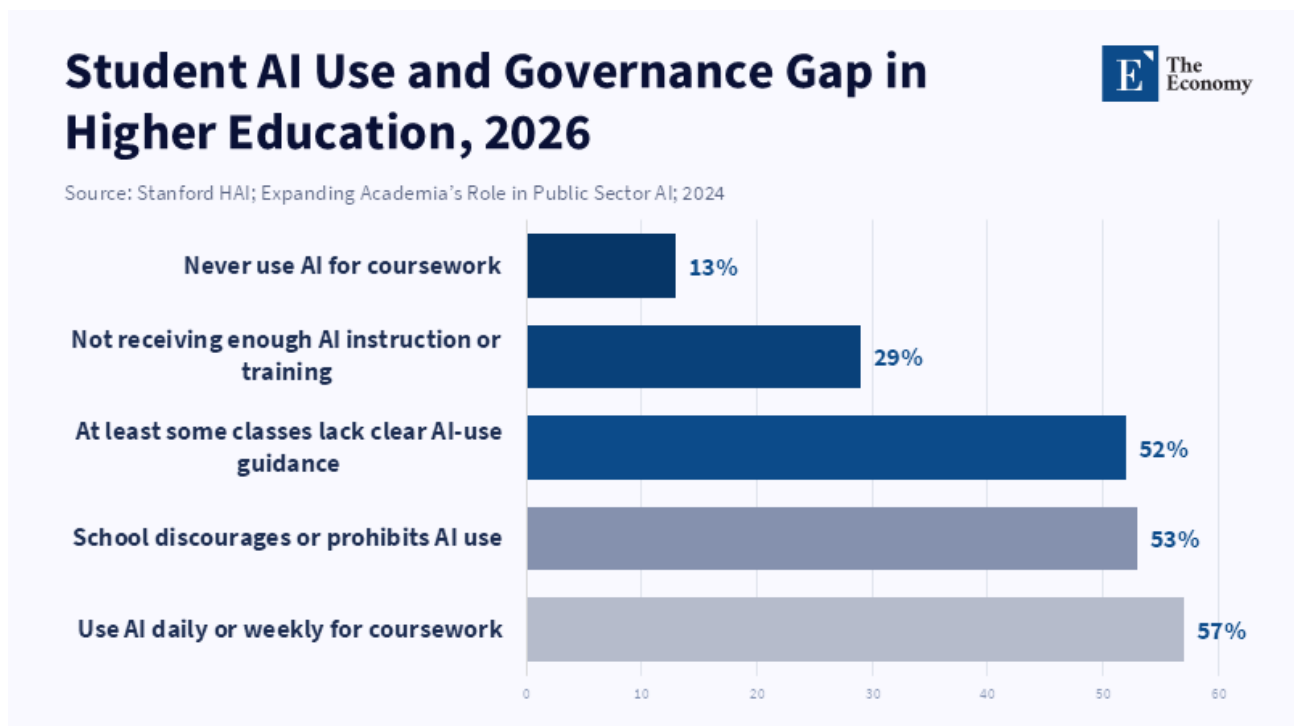


Figure 8: Universities already face mass AI use, but governance still trails practice.

It is also important for us to counter some strong attacks in the making effectively. One oft-cited defense is that AI is neither new nor revolutionary; in the long run, it will be no more disruptive in education than other technologies, such as calculators and search engines. But an accumulating body of evidence shows that treating generative AI as a calculator is "fundamentally misleading" (An AI education researcher)—it doesn't merely execute the standard mathematical operations; the machines are objectively superior to humanity at all that

they do (it takes less time and better) and it threatens to eventually take the place of the human (for writing essays and conducting research) completely. By consistently answering the homework questions and providing an information synthesis that seems to be a kind of shortcut to deep knowledge, they threaten to erode the human knowledge base that ought to be the very foundation of the calculator.

However, an argument can be made that AI will democratise knowledge ("opinion needs to be simplified and condensed to a form that is universally comprehensible and there's no other way of doing that than by reducing the integrity of that knowledge" (Edge 2023)), that it will be empowering for learners around the world. Certainly, open source AI tools represent this potentially global efficacy, but the data points to a more mixed picture.^[34] Yes, getting to the answer may seem easier than ever, but more shallow than ever for many students; yes, there are risks that generative AI can encode whatever biases and control structures its creators adhered to, if a handful of corporations cater the made 'answer' to the demands of what their (global?) customers want, they are merely determining terms and framing the answer. There are serious epistemic risks here that some topics will be playfully downplayed or more often upplayed in a corporate-centric presentation, with implications for wider democracy. The key question is whether the evidence for the positive far outweighs the evidence for its downside: Stanford (critical thinking), OECD (inequality), Stanford (journalism, displacing academics?).^[35] The key point is that the evidence for the benefits of "business-style AI labs" must be balanced against the evidence for their downsides.

Ultimately, universities cannot keep up with corporate AI labs and they cannot entirely give up on them, either. A more tenable, moderate middle ground must be found. Universities' things to do in an era of AI are to promote adaptive thinking, to do science ethically and for the benefit of the citizens who fund and use it and not to compete with private AI firms at major model training. Policy must step in to accommodate these shifts. Governments and funders might have to divert resources away from some expensive hardware, focused research to interdisciplinary projects or AI education. Accreditation agencies might have to set standards of AI literacy and ethics so both teachers and students can keep pace. Universities must solve the problems of so many students entering the workforce and seeking to learn more about AI through programs like internships or codesigned curricula and considering how to safeguard the integrity of the open scientific enterprise. These new roles must be taken seriously in policy, so the bright side of AI and higher education comes into view.

6 Conclusion - Redefining the University's Role in the Age of AI

The flow of AI talent from the academy to industry is not just a short-lived career phenomenon but a substantive change that affects education and society at large. From our work, it is clear that by the mid-2020s, the majority of the leading AI researchers are working in the private sector with an order of magnitude more compute and research funding to accelerate innovation than the academic researchers that inspired those advances. This trend is already slowly draining AI research from the academy, evidenced by a drop of around 1000 research papers a year, while AI patenting activity has increased.^[36] We also found that the infiltration of AI in students' learning and teaching is becoming problematic, with students using AI with no supervision and teachers struggling to keep their syllabi and assessments up-to-date.^[37] The report concludes that these trends have an effect on

humans' and institutions' cognition and gave an example of how a shallower learning effect could arise if AI does most of the "thinking" work and there is already a 10% decline in ads for non-routine jobs due to the penetration of AI in those domains.^[38] In addition, not all humans benefit from AI, with only those working in the top few percent of the workforce having the skills for AI. Interestingly, they tend to be in the highest paying, which suggests existing AI could contribute to a rise in inequality.^[39]

These are not inconsequential empirical findings; their implications should have us quickly dismissing naive equivalences between 'new' AI and given productivity gains (calculator-like interfaces, let's say, or new search engines). Such naive equivalents are clearly mistaken because generative AI fundamentally changes the way knowledge is produced and who has control over that process. Having charted how this change might shift the landscape of human capital and productivity, our argument should prompt a rethinking of new pathways for education and policy. For educators, the immediate takeaway is that educators will have more than content to teach; they will need to implement and cultivate sustainable cognitive skills and adopt AI only when it is a tool for understanding rather than replacing understanding with machine intelligence. This will mean rethinking curricula, adopting a pedagogy of conceptual mastery and bringing in evaluations that rely on a student's original argument or response rather than a student's eye or memorization of an AI's output. Administrators will need to add to these measures carefully defined policies on appropriate use, professional development on new pedagogies through which to administer any new curricula, academic integrity policies to create and enforce honor codes and support structures and detection services for both AI and human originality. For policymakers, the broader implications are far more far-reaching. Strategies for emerging human capital should account for the direction of AI tools and carbon, labor shifts that will follow; ministries should allocate funds to allow workers' reskilling and upskilling for AI, augmented careers; governments might consider investing in research facilities to partner with academia on AI safety and ethics; policies guiding new AI development should emphasize innovation that is safe and equitable—open source data sets, transparent models.

In conclusion, the report demonstrates that the AI revolution entails significant trade-offs. While AI delivers productivity improvements in the form of more efficient workflows and new tools in the short run, these depend on deeper transformations of incentives, human resources, and knowledge institutions. Countering these trade-offs requires research-driven approaches to policy: a rigorous effort to teach every learner about AI (per UNESCO), causal evidence to explore what works in edtech, creation of a market, and enabling AI flow through public investment. If not countered through deliberate action in the policy arena, gains will be concentrated among the likes of Google and Facebook while human skills gaps remain. As such, policymakers and educators would do well to heed the evidence we have assembled here and begin to adapt education policy, institutional structure and government strategy today instead of defaulting to the conventions of the past.

References

[14, 21, 27, 36] Akcigit, U., Chikis, C.A., Dinlersoz, E. and Goldschlag, N. (2026) *Attention (And Money) Is All You Need: Why Universities Are Struggling to Keep AI Talent*. NBER Working Paper No. 34964.

- [3, 5, 12] Akcigit, U., Chikis, C.A., Dinlersoz, E. and Goldschlag, N. (2026) ‘The great AI talent migration: Why universities are losing the future of innovation’, *VoxEU Column*, 11 April.
- [1, 17, 30] Amodei, D. (2024) ‘The Billion-Dollar Price Tag of Building AI’, *Time*, 15 April.
- [9, 11, 29, 35, 39] *Artificial Intelligence and Wage Inequality* (n.d.) OECD.
- [7, 24] ‘The Impact of AI Writing Assistants on Academic Writing Performance’ (2023) *Computers & Education*, 185.
- [6, 13, 16, 31] *The 2024 AI Index Report* (2024) Stanford Institute for Human-Centered Artificial Intelligence.
- [20, 23, 33, 37] Gallup (2026) *AI in Higher Education: Widespread Use, Unclear Rules*. Lumina Foundation, 2 April.
- [28, 32, 34] Goodier (2025) ‘Artificial Intelligence in Higher Education: A Global Statistical Synthesis for Policy and Quality Assurance Reform’, *Education Sciences*, 16(3).
- [26] HAI, S. (2025) ‘Economy’, in *The 2025 AI Index Report*. Stanford Institute for Human-Centered Artificial Intelligence.
- [2, 19, 25] HAI, S. (2024) ‘Research and Development’, in *The 2024 AI Index Report*. Stanford Institute for Human-Centered Artificial Intelligence.
- [38] ‘AI’s impact on the job market is starting to show up in the data’ (2026) *Axios*, 7 April.
- [8, 10, 22] Liu, Y., Yu, S. and Wang, H. (2025) *Labor Demand in the Age of Generative AI: Early Evidence from the U.S. Job Posting Data*. World Bank Policy Research Working Paper.
- [15] Mowreader, A. (2025) ‘AI Skills Needed in Many Postgrad Careers—Not Just Tech’, *Inside Higher Ed*, 1 August.
- [4, 18] Nolan, B. (2025) ‘Meta’s \$100 million signing bonuses for OpenAI staff are just the latest sign of extreme AI talent war’, *Fortune*, 18 June.