5 NOVEMBER 2026

Commentaries

New Series

AI LABS ECONOMICS

OR THE PRIVATISATION OF ARTIFICIAL INTELLIGENCE ECONOMIC RESEARCH

Simone Vannuccini*

COMMENTARY
NO.061NS/2026



*Chair of Economics of Artificial Intelligence and Innovation - Université Côte d'Azur/GREDEG CNRS

Senior non-resident Fellow at Fondazione CSF

The opinions expressed here do not necessarily represent the Fondazione CSF

Economic research is an important input to policy decisions. As policymakers devote increasing attention to artificial intelligence (AI), it should come as no surprise that the economics of AI has gained prominence in policy circles and within the economics profession. What is perhaps more surprising is that AI is also structurally reshaping the practice and organisation of economics: not only as a tool supporting analysis, but also through the entry of major AI industry players into the production of economic research itself.

In this commentary, I focus on this latter, ongoing trend: the rise of private economics research units established, funded and hosted by Big Tech companies and AI labs. Understanding this phenomenon is important for identifying the intellectual dependencies shaping policy-relevant economic research on AI. My primary interest is not what drives individual economists to leave academia and work for *Big AI* labs. Rather, I am interested in why AI companies have decided to enter the field of academic economic research, and what the implications of this move might be

But let's first take a step back to gain some perspective. AI is changing the rules of the game in science: from research design to paper preparation, from the publication pipeline to the development of early-career researchers' skills, many disciplines are experiencing a period of turbulence that could lead to drastic adjustments in how research is conducted and organised. The very *currency* of science, the journal publication, is facing a credibility crisis as AI-generated manuscripts increasingly flood submission systems and strain peer-review processes. At the same time, AI is accelerating scientific discovery and enabling previously prohibitive measurements and analyses. As research becomes increasingly automated, we risk an acritical reliance on large language models as modern oracles, and we also witness a trade-off between speed and collaborative work, with single-authored studies now on the rise at the expense of multi-authored ones.

Against this broad picture, one particularly important transformation is a shift in the organisations where scientific knowledge is produced, suggesting a growing *privatisation of AI research*. Here, we should distinguish between two related but distinct phenomena. The first is the outflow, or *brain drain*, of academic researchers working on AI from public universities to corporate research laboratories. The second is the expansion of these laboratories, traditionally focused on the technical aspects of AI, into economic research, increasingly producing their own analyses of the technology's broader economic consequences.

The *brain drain* from universities to private companies is certainly not a new phenomenon. Throughout modern history, researchers across different scientific fields have found productive havens for conducting

fundamental research within corporate laboratories – AT&T's Bell Labs in the United States being perhaps the most celebrated example, as Jon Gertner recounts in his 2012 book *The Idea Factory: Bell Labs and the Great Age of American Innovation*. Supported by AT&T's telecommunications business, Bell Labs gave us radical innovations in semiconductors (first and foremost, the transistor), satellite and wireless technology, programming languages and computing. It was also where Claude Shannon pioneered information theory – the very same Claude whose name was adopted by Anthropic for their family of AI models.

In the specific domain of AI, scholars have pointed to "the growing influence of industry in AI research" in recent years. AI research has always coexisted in the public and private sectors, but the field is now experiencing a structural change: not only are increasing numbers of AI doctoral graduates being hired by industry – an expected feature of a sector offering attractive employment opportunities – but professors are also leaving universities to lead research at frontier AI labs. This is potentially a concerning trend. The migration of established academics, with their research agendas, networks and intellectual leadership, may signal a more fundamental shift in where scientific knowledge is produced. At the same time, AI and computer science conferences are witnessing a growing presence of research papers authored by scientists working at AI companies. This development is consistent with Cecilia Rikap's argument that Google more than other Big Tech companies adopted a *university-like* posture to position itself as gatekeeper at the earliest stages of knowledge production, strengthening its ability to capture value across the whole value chain in markets relying on knowledge-intensive goods, services and technologies.

Until recently, the movement of AI researchers into industry was limited to computer scientists and other STEM researchers, alongside ethicists and philosophers working on questions of AI alignment. This predominance of technical expertise remains particularly visible in fields such as embodied or physical AI. Research on AI and robotics conducted by my own team has identified a dynamic we dubbed *Bell Labs 2.0*, with Google (and, to a lesser extent, Nvidia) acting as fundamental science laboratories and clearinghouses for scientific data and research findings. This finding points to a broader structural transformation: the possibility that the production and circulation of frontier scientific knowledge are increasingly led by privately controlled research ecosystems.

Beyond the question of where fundamental technical knowledge in AI is produced, a novel – and rather underreported – development is that AI labs are increasingly recruiting academic economists and establishing dedicated economics research units. Soumaya Keynes, writing in the *Financial Times*, was among the first to capture this trend in the making. Of course, economists have never worked exclusively in universities: many are employed by think tanks or NGOs, and corporate economic research itself has a long history. Bell Labs (again), for instance, made important contributions to operations research and industrial economics, notably to contestable markets theory, while AT&T's Bell System sponsored the *Bell Journal of Economics*, which eventually became the *RAND Journal of Economics*, one of the leading outlets for economic research. In Europe, and closer to home, Fiat maintained an important "office for economic studies" throughout much of the twentieth century, whose activities extended well beyond analysing trends in the automotive industry to encompass industrial policy, international economic affairs and European political economy.

The opinions expressed here do not necessarily represent the Fondazione CSF

Fondazione CSF

The opinions expressed here do not necessarily represent the Fondazione CSF

More recently, the rise of digital platforms during the 2000s and 2010s produced another wave of economists moving into the corporate sector, so much so that analysts detected the emergence of the *Techonomist in the Machine*. Platforms and tech companies recruited a particular breed of economists – those working in mechanism design, auction theory and behavioural economics – to help design markets (in particular, optimise pricing and product offers), study consumer behaviour through A/B testing and behavioural interventions (nudges), and analyse large-scale proprietary transactional data. Many economists welcomed the opportunity to finally put their abstract analytical models or laboratory experiment settings to work in real-world markets: beyond shiny headquarters, higher salaries and almost limitless resources, working inside tech companies allowed them to "hold a demand curve in [their] own hands".

The emerging generation of economics research units within AI labs appears to serve a somewhat different purpose. AI companies (some of which are also digital platforms) are now recruiting economists specialising in labour markets, productivity and economic growth. Rather than focusing on applying theoretical frameworks to fine-tuning their own market mechanisms and business strategies, these new research programmes are more empirical and outward-looking, investigating patterns of AI adoption, the future of work and broader economic scenarios. OpenAI, for instance, has launched its *Signals* programme to analyse real-world AI usage data, alongside an in-house Economists in Residence and an academic exchange programme. Google has engaged in a similar exercise with the establishment of an *AI & Economy Research Program*, drawing on its ATLAS study to investigate AI's implications for productivity, employment, skills and scientific innovation, while DeepMind has separately produced research on economic policy for AGI. Microsoft has formed an *AI Economy Institute*, structured as a think tank bringing together external scholars to study the societal transformation enabled by the "AI economy". Anthropic, meanwhile, has established a dedicated economics team whose Economic Index uses granular Claude interaction data to measure AI adoption and occupational exposure, informing also their macroeconomic scenarios for growth and employment. Beyond AI labs, even organisations traditionally associated with longtermist and x-risks programmes, such as the Future of Life Institute and the Windfall Trust, have recently turned *short-termist*, devoting attention to more immediate economic questions. The *economics-turn* in AI thus appears to be spreading across the whole AI research ecosystem.

As anticipated, the incentives for AI companies to play economist may be rather specific to the industry – although there are parallels with Bell Labs, whose economic research also contributed to debates over the legitimacy and organisation of the telecommunications sector. One can interpret AI labs' entry into economics as an attempt to shape the prevailing narrative about the benefits of their technologies. From this angle, corporate economic research clearly looks as one of the tricks in the "loose assemblage of techno-informational modes of influence that tech platform companies rely on to shape laws and regulatory proceedings" – what Elettra Bietti and Anne Bellon describe as *technological capture*. By producing economic evidence, AI labs can strengthen the public and policy-relevant legitimacy of their businesses, while influencing which questions are investigated, which data are made available and which interpretations gain prominence. Moreover, through empirical studies and scenario analysis, they can offer policymakers apparently rigorous tools for understanding AI's economic consequences, carrying greater authority than the more

speculative forecasting exercises produced by rather opaque initiatives such as AI-2027 or Europe2031.

The problem with this dynamic is that, as economic research produced within private laboratories gains legitimacy, academic research from universities may face a corresponding process of delegitimisation. The controversy surrounding *The Economist*'s attack to Daron Acemoglu – interpreted by some as a proper proxy-war, supported by the AI labs, against a prominent critical voice – is the best illustration of this tension. Acemoglu (together with his co-authors) has dared to argue that AI should be "*pro-workers*", that its macroeconomic contribution to productivity growth may fall considerably short of expectations even in the rosiest scenarios, and that automation technology may be used to support repression and undermine democracy. Now that AI labs have their own *cool* economists, they may be in an increasingly strong position to undermine critical arguments coming from *less cool* (and less data-endowed) academic researchers from the public sector.

It goes without saying that, this time around, the opportunity to (paraphrasing) "*hold a production (or a labour demand) function in your hands*", thanks to the trove of data generated by AI conversations and, increasingly, agentic interactions, may be too tempting for many economists to pass up. Yet beyond these individual advantages, the implications of the privatisation of AI economic research for the economics profession and the public interest may outweigh the benefits of exploiting exciting and novel datasets. The conditions under which AI economic knowledge is produced and disseminated will depend on the interests of the companies financing it; economists working in corporate AI laboratories may find their research agendas narrowed and constrained by their employers' priorities, while also facing ethical dilemmas over the disclosure of findings that do not align with the prevailing narrative of AI as an overwhelmingly positive and transformative technology.

AI labs are not inventing corporate economic research, but potentially reviving and transforming an older institutional arrangement in which powerful technology producers become influential sources of economic knowledge about their own industries and their wider social consequences. Given this trend and its inherent risks, how can institutions supporting research in the public interest respond? The trivial-to-see-but-hard-to-implement answer is to make academia more attractive, so as to limit the brain drain of economists and counter technological capture. This would require not only better salaries and career opportunities, but also greater access to research funding, data and computational resources. While few would disagree in principle, such an effort is difficult to envisage in times of tight public budgets, limited political will and widening inequalities in means between public and private actors. Governments are also increasingly prioritising alternative investment directions, particularly defence spending, because of growing geopolitical fragmentation. Defence-funded research can certainly generate civilian technological spillovers and end up feeding back economics, as it has done historically (*e.g.*, in the case of operation research), but relying on such indirect benefits is a rather tortuous alternative to financing universities directly.

A different answer is for the public sector to imitate (and improve upon) the initiatives of the private sector. A tit-for-tat response, if you like. To some extent, this is what is happening in the United Kingdom with the recently

The opinions expressed here do not necessarily represent the Fondazione CSF

Fondazione CSF

launched AI Economics Institute, which adopts elements of the AI labs' economics research model within a public institution. While such an institutional innovation is by no means immune to technological capture (on the contrary, it may be particularly vulnerable to it), and it also requires significant resource allocation, it represents an interesting attempt to level the playing field in economic research on the impacts of AI. The European Union should take note, developing its own AI economics research unit capable of competing with those of Big AI, while building strong collaborations with universities and the many independent research institutes already active across the World that pursue the mission of studying AI in a transparent, critical way.

The opinions expressed here do not necessarily represent the Fondazione CSF

Fondazione CSF