

AI & Market Power

Findings from
Fellows

June 2026

European
**Artificial Intelligence
& Society Fund**

CONTENTS

Executive Summary	3
Section 1: Big Tech’s stranglehold on AI markets	6
Section 2: Big Tech’s grip on non-tech sectors	12
Section 3: The Fellows	16

Executive Summary

The AI & Market Power Fellowships were supported by the [Open Society Foundations](#) and administered by the [Institute of International Education](#). The European AI & Society Fund was responsible for the selection of the fellowships and for hosting the fellowships community.

In November 2024, the European AI & Society Fund backed eight Global AI & Market Power Fellowship teams, made up of 14 individuals, to spearhead investigations into market power dynamics in the AI ecosystem.

- The Fellows investigated how large US technology firms dominate digital infrastructures, increasing dependencies by AI startups, SMEs and governments on their services. Their research demonstrates that concentration in AI markets is not a future risk, but a deepening bind with every investment round, cloud credit programme, and model-infrastructure bundle.
- The same three US companies – Microsoft, Amazon and Google – control the cloud layer, dominate model distribution, shape startup ecosystems through venture investment, and are extending their reach into adjacent sectors from energy to biology to defence.
- The window for intervention is narrowing: lock-in mechanisms compound over time, the evidence base needed to act is actively shrinking, and “sovereign AI” initiatives risk channelling public money into structures that entrench rather than challenge existing market power. This requires action from competition authorities and heads of state alike, but the nature of that action differs.
- Artificial Intelligence should benefit everyone, not just the powerful few. But we live in a world where the direction of AI has been captured by a small number of tech giants, set on amassing vast wealth and entrenching their power. Trapped in their narrow vision for the future, we are failing to harness the potential of AI technologies to build the fair, democratic and sustainable societies we all want to live in.

METHODOLOGY

This report was prepared by Frederike Kaltheuner and the European AI & Society Fund. It is based on three categories of source materials: (a) funding/ narrative reports submitted by grantees to the Institute of International Education (IIE) or European AI & Society Fund, (b) the actual research outputs produced with fellowship funds, and (c) notes from joint workshops with the fellows in February and December 2025. At the time of writing this analysis (June 2026) not all projects had been published or finalised, and some of the reporting was not available.

AI & Market Power Fellows and their research topics – more detail about each Fellow and their findings can be found in Section 3:

Project 1

DAVID GRAY WIDDER

Forecasting Cloud Capture:
Tracking Cloud Giant
“Investment” in AI Startups

Project 2

FRISO BOSTOEN AND GIORGIO MONTI

Fair and Inclusive Markets for AI
Training Data

Project 3

ISHA SURI, SHYAM KRISHNAKUMAR AND TITIKSHA VASHIST

Democratising AI Infrastructure
with Open Cloud Compute

Project 5

JIM THOMAS

The ‘Generative Biology’ Grab
– Tech Titans Takeover the
Bioeconomy

Project 5

LOUIS BOYD-MADSEN, GABRIELE DI DONFRANCESCO AND CLÉMENT POURÉ

Mapping the uptake and
effects of algorithmic trading in
Europe’s energy market

Project 6

MARGARIDA SILVA

Value chain of genAI startups

Project 7

NICOLAS ESCHENBAUM

Market Power and
dependencies in the Market for
Foundation Models: an Empirical
Investigation

Project 8

YUNG AU AND SRUJANA KATTA

The political economy of military
general-purpose AI systems

Big Tech's stranglehold on AI markets

Five fellowship projects collectively mapped concentration in the commercial AI stack: cloud infrastructure (Widder, Suri), foundation models (Eschenbaum), training data markets (Bostoen and Monti), and the value chain from chips to distribution (Silva).

1.1

Cloud concentration is the foundational layer

Competition authorities have focused on challenging large individual partnerships (Microsoft-OpenAI, Amazon-Anthropic). But fellowship research shows this misses the main picture.

Widder finds that Big Cloud's market-shaping power comes from thousands of smaller investments that cumulatively create financial dependence across the startup ecosystem.

Suri's interviews with Indian AI-startup companies found they are locked into using cloud storage with Amazon Web Services (AWS), Google Cloud Platform (GCP), and Microsoft Azure through free credits, VC-hyperscaler partnerships, and bundled offerings. These small companies are steered toward hyperscalers one decision at a time, with virtually no viable alternatives.

Corroborating this dependency, Silva finds 10 of 12 top genAI startups rely on the same three hyperscalers for infrastructure and nine commercialise via their platforms.

The cumulative effect is what matters, and current scrutiny frameworks are not designed to capture it. In our December workshop, Fellows reflected that market power regulation should "focus on cloud even in non-cloud projects".

1.2

Vertical integration is more complex than often assumed

Multiple fellowship projects show that Big Cloud companies are integrating across the AI stack in ways that escape traditional scrutiny.

Widder finds that thousands of small venture investments create "soft" vertical integration: Big Cloud invests in data companies, infrastructure-as-a-service providers, and developer tools, building financial dependence that compounds technical and contractual lock-in.

1.1

"Market power regulation should focus on cloud even in non-cloud projects."

Silva's value chain mapping confirms this from the supply side: the same companies that control cloud infrastructure also control the platforms where genAI models are commercialised.

Suri shows this extending to the foundation model layer, where proprietary models are increasingly bundled with cloud offerings, making each layer harder to contest independently. The resulting lock-in operates through multiple reinforcing mechanisms simultaneously: credit subsidies, pricing opacity, labour market standardisation around hyperscaler tools, contractual minimum-spend obligations, and marketplace distribution.

This is not a series of separate competition concerns but a single integrated strategy.

The vertical integration documented across these projects implies that behavioural remedies (which focus on modifying a company's conduct rather than its structure) are unlikely to be sufficient on their own. When the same company controls cloud infrastructure, model distribution, and startup financing simultaneously, addressing any single layer leaves the others intact. Contractual, technical, and financial lock-in reinforce each other.

The Fellows' consensus was that behavioural remedies are failing, and that more active market shaping is the path forward.

1.3

Market dynamics can differ significantly across sectors and use cases

A blanket definition of "foundation models" as a single market will systematically underestimate concentration where it is most entrenched and overstate competition where it is most visible.

Eschenbaum's findings show Big Tech has the monopoly on AI in healthcare, the oligopoly on AI used for coding, and a moderate chunk of the content generation market. Meanwhile Boyd-Madsen et al. find analogous variation of AI use in energy markets.

These have direct implications for how competition authorities define relevant markets.

Authorities investigating AI markets need to disaggregate by sector and use case, or risk drawing conclusions that are technically correct in aggregate but wrong everywhere that matters.

1.3

Authorities investigating AI markets need to disaggregate by sector and use case, or risk drawing conclusions that are technically correct in aggregate but wrong everywhere that matters.

1.4

Competition authorities lack adequate legal tools and are not using those they have

Fellowship research shows that data workers and content creators bear significant costs from AI market concentration: precarious conditions, suppressed bargaining power, and value extraction by intermediaries. Yet competition authorities currently ignore labour concerns, despite them reflecting market power.

Widder shows that deal-by-deal scrutiny misses cumulative investment patterns, and Eschenbaum's vertical stratification finding implies that horizontal market definitions are the wrong frame.

Silva's Meta-Scale AI case study tested whether antitrust authorities would investigate labour concerns and found they considered them outside their scope.

Suri notes that India's Digital Competition Bill, which would have categorised cloud as a core digital service, was withdrawn.

Bostoen and Monti's work on training data pricing makes the connection explicit: without creating markets that ensure compensation for creators, the value of their work is captured by model developers with dominant market positions.

Competition frameworks that exclude labour from their scope will systematically miss a significant dimension of how market power operates in AI.

Fellows agreed that market shaping (including structural separation and utility-style regulation, or an open market for computing) deserve serious consideration to address Big Tech's AI market concentration.

1.5

Opacity is pervasive and access to evidence is at risk

The fellowships document a paradox: the more concentrated AI markets become, the harder it gets to produce the evidence needed to challenge that concentration.

Datasets are being paywalled, model cards are disclosing less, pricing remains opaque, and organisations are reluctant to admit lock-in. If policymakers want evidence-based regulation of AI markets, they need to ensure the evidence exists. This could mean mandatory disclosure requirements for cloud pricing, investment portfolios, and model training data; preserving public access to commercial datasets; and funding sustained monitoring infrastructure rather than relying on one-off research.

1.4

Competition frameworks that exclude labour from their scope will systematically miss a significant dimension of how market power operates in AI.

For example, Silva's work relied on model cards – which Google defines as “simple, structured overviews of how an advanced AI model was designed and evaluated” – that already disclose less now than two years ago.

Widder's Crunchbase dataset – which provides data such as company name, type, social media links, number of investors, and much more – was paywalled mid-project, now costing thousands of dollars.

Access to sources is also challenging. Suri describes a “culture of silence” that only breaks in confidential interviews. And Bostoen and Monti note the training data market is “very opaque” and that access to non-disclosed agreements would transform the field. Eschenbaum's sample of 19 reflects the difficulty of getting organisations to admit lock-in to Big Tech contracts.

This creates a paradox: the more concentrated AI markets become, the harder it gets to produce the evidence needed to challenge that concentration.

The evidence being collected by the fellowship is new and often ahead of where regulators and the public debate currently sit, but the window for producing it may be closing.

1.6

“Sovereign AI” framings miss market concentration dynamics

The growing political appetite for “sovereign AI” risks channeling public money into initiatives that leave the underlying market concentration untouched or further entrenched. National debates around “sovereign AI” and “sovereign cloud” in Europe, India, and elsewhere largely fail to engage with the market concentration that makes sovereignty difficult in the first place.

Suri finds that most sovereign stack proposals in India operate only at the application layer while the cloud, model, and data layers remain controlled by three US companies – Amazon, Google and Microsoft.

Silva's finding that the European Parliament's AI model choice was constrained by a pre-existing cloud procurement agreement with Amazon illustrates how lock-in renders sovereignty aspirations moot in practice.

Widder's analysis of Big Cloud's accelerator strategy abroad shows how “sovereign AI” investments can function as tools of market expansion by the very players they claim to counterbalance.

The sovereignty debate tends to fall into two traps: either assuming that layer-by-layer technical solutions are sufficient without addressing the underlying market structure, or seeking to replicate the US model with

1.6

The growing political appetite for “sovereign AI” risks channeling public money into initiatives that leave the underlying market concentration untouched.

national monopolies, which, as Suri's research concludes, reproduces the same harms under a different flag.

Sovereignty policy that does not address vertical integration, switching costs, and lock-in at every layer of the stack is sovereignty in name only. Equally, replacing foreign monopolists with domestic ones reproduces the same harms. Public investment in AI infrastructure should come with structural conditions: interoperability, portability, and limits on vertical consolidation.

1.7

The need to articulate an affirmative vision for AI

The fellowships were deliberately focused on generating evidence about AI market power. But a recurring theme is the gap between the sophistication of the research and where public debate and regulatory understanding currently sit. The gap is not only one of communication but of narrative: public debate remains anchored in framings (e.g. "no abuse of dominance," "AI national champions") that the fellowship research has empirically dismantled. Translation work therefore requires not just simplification but active narrative challenge.

Although beyond the remit of these fellowships, future grants could explicitly fund translation activities: short policy briefs, journalist briefings, civil society-facing visuals, podcast appearances. Some fellows will be natural communicators; others will need institutional support or co-authors with translation skills. Budgeting for a communications partner or for professional editing and design (as Au and Katta did with Open Tech Fund money for illustrations and typesetting) could be an option.

The cohort concluded that the critique of concentration is well-developed but the affirmative vision is not. The final workshop captured this bluntly: "Pushback to concentration argument is: so what is the alternative?" Ideas were raised (cloud as a utility, structural separation, open compute marketplaces, collective bargaining for content creators) but none yet has the traction the critique does.

Future programmes could explicitly fund solution-oriented work: scenario exercises, comparative policy analysis of structural remedies, and engagement with actors already building alternatives. Several fellows noted that "sovereign stack" proposals typically only operate at the application layer while the lower layers remain foreign-controlled, suggesting that even the most prominent alternative framings need critical scrutiny.

Big Tech's AI market power is clear but the alternative is not. More needs to be done to articulate an affirmative vision for AI that serves people, society, and planet.

1.7

Big Tech's AI market power is clear but the alternative is not.

More needs to be done to articulate an affirmative vision for AI that serves people, society, and planet.

2

Big Tech's grip on non-tech sectors

2

Three projects focused on how market concentration in AI affects adjacent markets and domains: energy (Boyd-Madsen, Donfrancesco and Pouré), biology (Thomas) and the military (Au and Katta).

2.1

AI market power reproduces itself in adjacent sectors

All three projects trace how the concentrated Big Tech-owned AI stack extends into domains with their own pre-existing power asymmetries.

Thomas shows Big Tech moving into genetic engineering via “generative biology,” a domain that has long been characterised by biopiracy and the exploitation of indigenous communities.

Boyd-Madsen et al. show algorithmic trading software vendors becoming a new intermediary layer in European energy markets, reinforcing the market position of incumbent energy majors.

Au and Katta show how commercial AI firms gain legitimacy, revenue, and data access through military contracts.

In each case, the AI stack does not enter a neutral market; it layers on top of existing inequalities and amplifies them.

2.2

The state underwrites the market expansion and potential AI bubble burst

A shared finding is that these sectors are not purely commercial.

Au and Katta reflected on how military AI is directly state-funded through procurement.

Thomas identified that generative biology benefits from state research funding, biodiversity agreements, and “sovereign” framing that channels public money toward private biotech.

Boyd-Madsen et al. outlined how AI concentration in energy markets is shaped by regulation, subsidies, and the politics of the clean-energy transition.

2.2

“The market does not decide itself when its viability depends on government money and political backing.”

The final Fellows workshop captured this bluntly: “the market does not decide itself when its viability depends on government money and political backing.”

The fellowship research on these adjacent markets shows that state procurement, subsidies, and security-framed investment increasingly underwrite the AI stack. If the consumer AI bubble bursts, public money may be what sustains the market.

Governments should not treat the AI market as if it decides itself when its viability depends on state support. Public investment in AI – whether through military procurement, sovereign cloud initiatives, research funding, or industrial strategy – should come with conditions to ensure they serve the public interest and do not further entrench concentration. The alternative is that taxpayers fund the entrenchment of the very market power that sovereignty initiatives claim to counter.

2.3

Extraction from communities that cannot bargain

Each project documents value flowing from vulnerable or structurally disadvantaged actors toward already-powerful ones.

Thomas documents communities and parks worldwide providing genetic material to companies like Basecamp Research, with little value returning; while Au and Katta document data extracted in conflict zones.

Boyd-Madsen et al. show that “flexibility rents” from algorithmic trading ultimately come at the expense of consumers, small renewables owners, and less tech-savvy competitors.

The common pattern is extraction from actors who lack the information, leverage, or institutional standing to negotiate terms. More needs to be done to protect communities, and shift the debate from corporate or national self-interest, to justice, ecology and public interest.

2.4

Regulatory frameworks lag behind or are actively shaped by incumbents

Boyd-Madsen et al. show that Regulation (EU) 2024/1106 (REMIT II) – which came into effect in April 2026 and aims to significantly expand and intensify the EU wholesale energy market transparency regime – captures only the final stage of trading and asset-optimisation tools and has been weakened by energy trading industry lobbying.

Thomas finds that existing biodiversity governance (the Convention on Biological Diversity, the Nagoya Protocol) was not designed for AI-driven sequence data extraction and is struggling to keep pace.

2.3

The common pattern is extraction from actors who lack the information, leverage, or institutional standing to negotiate terms.

Au and Katta find military AI supply chains operate in a space where transparency norms barely apply and oversight is fragmented across jurisdictions.

In all three cases, the gap is not simply “regulation is slow” but that industry actors are actively involved in shaping or circumventing the rules.

2.5

Opacity is a feature, not a bug

All three sectors are structurally opaque, but for different reasons.

1. Military supply chains are opaque by design (national security classification, corporate confidentiality).
2. Energy-market algorithmic trading is opaque because the software layer sits between regulated entities and is not itself subject to disclosure requirements.
3. Generative biology is opaque because sequence-data deals, bioprospecting partnerships, and IP arrangements are commercially confidential and span jurisdictions with different disclosure norms.

Each project had to develop workarounds: FOIA requests and industry-conference attendance (Boyd-Madsen), triangulation of procurement records and open-source investigation (Au and Katta), movement networks and investigative journalism (Thomas).

2.6

These sectors legitimise the AI stack in ways consumer AI cannot

A point that emerges from the final Fellows workshop is that the “AI bubble” narrative applies most directly to consumer-facing generative AI.

Military AI, energy optimisation, and biotech/pharma are where the “meaningful AI” narratives live: national security, curing cancer, fighting climate change, keeping the lights on. These narratives provide political cover for continued investment even if consumer genAI disappoints.

This means the sectors covered by these three projects are not peripheral; they may be where the AI stack’s economic viability is ultimately secured.

2.6

Military AI, energy optimisation, and biotech are where the “meaningful AI” narratives live: national security, curing cancer, fighting climate change. These narratives provide political cover.

3

The Fellows

Project 1

Forecasting Cloud Capture: Tracking Cloud Giant ‘Investment’ in AI Startups

by David Gray Widder and Nathan Kim



Project 1 · Team of 2

Focus

Big Cloud investment

AI supply chain

AI startups

Fellow profiles

DAVID GRAY WIDDER studies epistemic and political economic power in AI production. He is a postdoctoral fellow at Cornell Tech, and earned his PhD from Carnegie Mellon University. He is an assistant professor in the School of Information at the University of Texas at Austin.

NATHAN KIM is a PhD student at the University of Michigan School of Information, and a research intern at the Distributed AI Research Institute. He is interested in critical studies of race, finance, logistics, and tech infrastructures.

Output

[How Big Cloud becomes Bigger: Scrutinizing Google, Microsoft, and Amazon’s investments ↗](#), D. Widder and N. Kim, 4 August 2025, SSRN.

Project purpose

Cloud giants are building their empires. Google, Microsoft, and Amazon seek to capture more of the AI cloud market, but what are their tools and strategies for empire building? How might we anticipate the state of their empires tomorrow? These Fellows analysed historical datasets of startup investments to track how cloud giants are investing in AI startups to extend their cloud market power, how this changes over time, and how they may do so in the future.

They set out to demonstrate a significant new competition concern: financial dependence through cloud conglomerate investment in the startup ecosystem and how it relates to past market concentration concerns in AI. The project challenged the prevailing regulatory approach that focuses on individual large deals (e.g. Microsoft-OpenAI), arguing that competition authorities should scrutinize the cumulative, ecosystem-wide effect of thousands of smaller investments that matters.

Key findings

- Big Cloud (Amazon, Google, Microsoft) invests as frequently and at similar scale to the largest VC firms, and about ten to a hundred times more than other Big Tech companies; the rest of the “Magnificent Seven” barely register by comparison.
- When Big Cloud invests in AI supply chain companies, it is disproportionately likely to be the sole or lead investor, creating financial dependence that compounds other lock-in mechanisms such as high egress fees and minimum-spend contracts (Sec. 2.4 of research paper)
- Accelerator programmes steer the market toward technologies beneficial to Big Cloud’s long-term interests, targeting startups likely to require significant cloud compute and promoting building around the investor’s offerings; abroad, these programmes serve as tools of market expansion, tethering recipient states’ innovation systems to US-based cloud infrastructure and billing relationships
- Rivals cannot examine the prices they must compete with, enabling incumbents to tailor discounts to prevent startups from switching
- In AI especially, this amounts to “soft” vertical integration that escapes regulator scrutiny

Key learnings

- Vertical integration is more insidious than previously understood; it takes rigorous academic research to document and understand how Big Cloud asserts its power
- Research question changed as the analysis progressed
- The Crunchbase dataset used for the analysis is no longer publicly available and now requires thousands of dollars to access, illustrating the fragility of the data infrastructure underpinning this kind of research

Reach

Impact activities included closed-door presentations to Dutch and UK competition authorities, interest from the German competition office, and engagement with former US FTC officials; the paper was flagged as a Top Papers List entry on SSRN. The paper also appeared in Tech Policy Press.

Fair and Inclusive Markets for AI Training Data by Friso Bostoen and Giorgio Monti



Project 2 · Team of 2

Focus

AI training data market

Intellectual property lawsuits

Content creator safeguarding

Fellow profiles

FRISO BOSTOEN is an assistant professor of competition law and digital regulation at Tilburg University. He researches how the law can deal with emerging technologies in a way that fosters innovation and fairness.

GIORGIO MONTI is a Professor of Competition Law at the Tilburg Institute for Law, Technology and Society (TILT). His research interests cover competition law, the regulation of digital markets and European Union Law.

Output

[Putting a Price on AI Training Data ↗](#), F. Bostoen and G. Monti, 28 November 2025, SSRN.

Project purpose

- Examines how markets for AI training data can be structured to be both inclusive for creators and fair in remuneration
- Addresses the most difficult question: how to put a price on AI training data
- Motivated by the observation that the training data market is emerging but highly opaque, with terms kept secret and developers like Meta willing to walk away rather than pay
- Notes that intellectual property lawsuits are not a systematic solution, and licensing deals have been limited to a few large publishers
- More comprehensive legal and technological regimes have been proposed, but they do not reckon with the pricing question

Key findings

- Surveys four existing regulatory frameworks for lessons on market design for price formation: Data Acts, press-publisher bargaining models, excessive pricing doctrine, and standard essential patent (SEP)
- The owners of AI training data that do not have exclusive access are in a weak starting position; even if access is prevented, a regime still needs to incentivise both parties to negotiate a price
- Any framework should be balanced toward creators but should not make AI development impossible
- Key lesson: price formation will not happen unless a scheme identifies the purposes for remuneration first and then creates incentives for the parties to come to the table and negotiate; this is usefully supplemented with a pricing model and safeguards to prevent the use of economic power to drive prices up or down
- Four different models have different market power mitigation measures; collective bargaining of content creators has the most potential
- Smaller publishers in France and Australia have been able to bargain with Google and Meta, though larger players like the New York Times may not want to collaborate in this way

Key learnings

- The market is starting to form but remains very opaque
- Content-creator lawsuits (including the Anthropic case) may be starting to affect data pricing
- Scope expanded beyond original proposal
- Access to non-disclosed agreements (e.g. via court disclosure) would significantly strengthen future work

Reach

Output published as TILEC Working Paper on SSRN; presented at “Competition Law 2.0” conference at UCL; target publication in GRUR International.

Project 3

Democratising AI Infrastructure with Open Cloud Compute by Isha Suri



Project 3 · Individual Project

Focus

AI startups

Hyperscaler dominance

Foundation model dependency

Fellow profile

ISHA SURI is a Research Lead at the Centre for Internet and Society, India. Her research interests include competition policy, digital infrastructure, and platform governance. She is an Electrical Engineer and also holds a post-graduate degree in law from the Indian Institute of Technology, Kharagpur.

Project purpose

Despite hyperscaler dominance (Google, Amazon, and Microsoft control nearly two-thirds of the global cloud market and over 85% in India), the dominant narrative in India remains that there is no “abuse of dominance”.

Based on 15 in-depth interviews, the research examines cloud concentration in India’s AI development layer, its impact on domestic providers, and whether a “Digital Public Infrastructure” approach could offer a viable alternative.

Key findings

- Choice and concentration: AWS, GCP, and Azure dominate across firm types; even “multi-cloud” setups keep usage within the hyperscaler ecosystem. Oracle and domestic providers are marginal
- Adoption driven by familiarity, not technical merit: CSP selection often follows founder/leadership alumni ties (“schmooze network” of IIT/IIM graduates) and engineer employability incentives, not technical superiority
- Credits as market steering: All startups interviewed built on free credits; one called this a “drug dealer model.” VC-hyperscaler partnerships create “soft coercion.” Significant pricing opacity, with discounts of 20-80% available only to firms with bargaining power
- Lock-in and barriers to exit: Switching is technically possible but economically unviable. In one case, a startup that cut GCP spending by 85% had GCP invoke minimum-spend obligations to block further migration. Proprietary services accumulate technical debt over time
- Foundation models as a new bottleneck: Models bundled with CSPs (Gemini + GCP, Azure + GPT, AWS + Claude). All respondents used proprietary models; no domestic companies developing indigenous ones
- Single point of failure: All respondents agreed an outage would be “catastrophic,” yet resilience plans do not account for it

Key learnings

- Utility framing vs. reality: Consensus that cloud is “essential infrastructure” comparable to electricity, but CSPs remain privately owned without public obligations. India’s Digital Competition Bill, which could have imposed such obligations, has been withdrawn
- Alternatives: Broad consensus on the need to reduce hyperscaler dependence, but uncertainty about what this entails in practice domestic CSPs (scepticism about scale), public compute (surveillance concerns), open compute grids resembling UPI (far harder than earlier DPI initiatives)
- Policy recommendations: Regulate CSPs as utilities; structurally separate AI supply chain layers; implement ex-ante competition laws

Project 4

Unmasking Generative Biology: Tech Titans take over the Bioeconomy by Jim Thomas



Project 4 · Individual Project

Focus

Generative biology

Biopiracy

AI training data

Fellow profile

JIM THOMAS of scanthehorizon.org is an international civil society researcher and advocate addressing new technologies, biodiversity, food systems and corporate power. He has three decades of experience tracking emerging technologies and international governance processes while supporting movements to understand tech impacts and craft response strategies.

Output

[AI's Large Looting Models? The Emerging Generative Biology Stack as the Next Frontier of Biopiracy ↗](#), Jim Thomas, ETC Group, 2024.

[Hoodwinked in the Hothouse](#), Substack [Scan the Horizon](#). Thomas was also quoted in [GPE News Docs](#), [The Financial Times](#), and interviewed on the [Breaking Green Podcast](#).

Project purpose

- Investigates how large tech players are moving into genetic engineering via “generative biology,” where AI models trained on DNA, protein, and other biological data generate novel proteins, DNA sequences, or whole organisms.
- Frames AI-driven collection and sequencing of biological data as a new frontier of biopiracy, where digital sequence information (DSI) on genetic resources is the “oil” of this emerging sector.
- Develops an initial taxonomy of where AI and genetic engineering are overlapping, surveys the entities and activities shaping the GenBio ecosystem, and presents a conceptual framework of the “GenBio stack”.

Key findings

- The GenBio “stack” is incredibly expensive and far from ready; its main value currently is generating good PR for generative AI and attracting government, military, and philanthropic funds that may prop up gen AI infrastructure after the bubble bursts.
- Market concentration follows patterns familiar from other AI sectors but with a twist: those who control the translational technologies bridging digital and biological domains (computation, sequencing, and synthesis) will hold disproportionate influence over the entire stack.
- GenBio converts common goods into private assets, enclosing biodiversity in ways that benefit investors while sidelining steward communities; at the forefront is Basecamp Research, paying Global South communities to collect biodiversity from their own territories with little value returning to them
- The sector is increasingly courting military and security institutions, leveraging anti-China sentiment, while the “black box” nature of generative models leaves regulators effectively blindfolded when approving products whose underlying logic is opaque

Key learnings

- The project pivoted significantly towards informing and supporting social movements and civil society groups to understand and address the threats of AI in their own terms, connecting food sovereignty, climate justice, and indigenous rights movements.
- Key entry points for civil society engagement include negotiations under the Convention on Biological Diversity, the FAO Seed Treaty, privacy and data protection enforcement, and national/regional policy on AI training data.

Reach

- Thomas presented his findings in Canada, the USA, Kenya, Ethiopia, Zambia, South Africa, India and France, helping educate and empower activists and civil society groups to understand and address the threats of AI and biotech.
- He also informed international strategy meeting on AI and climate justice for Sunrise Project and the Global Fund for a New Economy; organised a 3-day workshop on AI and Environmental Justice at the GFNE’s global convening in Kenya; and attended the first ever Mining and AI tech conference, on behalf of indigenous movements.
- In addition to publications, he co-designed and co-led a track on AI and Environmental Justice at the GNI Global Convening in Nairobi (September 2025), forming a cohort of allies whose common agenda has become the basis for the “EarthKeepers vs AI Empires” convening in Lusaka, Zambia (May 2026, targeting up to 150 activists from all continents).

Project 5

Mapping the uptake and effects of algorithmic trading in Europe's energy market

by Louis Boyd-Madsen, Gabriele Di Donfrancesco, and Clément Pouré



Project 5 · Team of 3

Focus

Algorithmic trading

Software intermediaries

Energy market manipulation

Fellow profiles

LOUIS BOYD-MADSEN is an investigative journalist focused on the socio-economic impacts of the tech, climate, and finance intersection. His previous work as a university lecturer at the Institute for American Universities examined the political economy of nationalism and globalisation.

GABRIELE DI DONFRANCESCO is an investigative journalist focused on utilities, climate, and social issues, often encompassing the impact of new technologies. Bylines on the Bulletin of the Atomic Scientists, Jacobin, and RomaToday, among others.

CLÉMENT POURÉ is an investigative journalist reporting on state surveillance and technology at work. His investigations have been published in Mediapart, Le Monde and the international French-speaking press

Output

[Comment l'IA détraque le marché européen de l'énergie ↗](#), Louis Boyd - Madsen, Gabriele Di Donfrancesco and Clément Pouré, Avril 2025, Reporterre.net.

[L'IA, nouveau joujou dangereux du marché de l'énergie ↗](#), Louis Boyd - Madsen, Gabriele Di Donfrancesco and Clément Pouré, Mars 2025, Le Pave numérique.

Project purpose

- Maps the software intermediaries now acting as a new layer of "algorithmic infrastructure" in the European energy market, which dictates how energy is generated and traded.
- Seeks to understand how these actors shape market power in the energy sector, how they constitute a new form of market power in and of themselves, how they are being regulated, and how those regulations have been shaped by lobbying.

Key findings

- The European energy market is increasingly mediated by a small set of software providers, with one specific intermediary holding contracts with the vast majority of European energy majors and scaling up via M&A deals with competitors.
- These technologies are very likely enabling existing large energy majors, with mixed portfolios of hydro, batteries, and gas, to extract substantial “flexibility rents” from the energy market by exploiting the volatility in intraday/day-ahead prices that has emerged as the energy system transitions to clean power
- Main takeaway: new intermediaries provide competitive advantage to legacy energy companies, creating new opportunities for fossil fuel value extraction through machine learning forecasting, automation, and high-frequency trading
- Renewables are more volatile, which limits how much companies can speculate on them; renewables should make energy cheaper for consumers, but software enables energy companies to keep good margins of profit at the expense of consumers
- REMIT II regulations in their current form ignore all but the final stage of these trading/asset optimisation tools and have been subject to heavy lobbying from the energy trading industry
- Regulators are not looking at concentration of the software intermediaries

Key learnings

- The project pivoted from an initial framing around algorithmic collusion to a focus on entrenched market power via exploitation of volatility
- The Dutch regulator appeared to be the only one paying attention to the role of intermediary players and algorithmic collusion, and was transparent and eager to engage
- Other regulators tried to slow down the investigation; in many cases they were aggressive, avoided answering, or gave scattered information
- Final output was being coordinated with media outlets for publication at time of reporting

Value chain of genAI startups

by Margarida Silva



Project 6 · Individual Project

Focus

GenAI startups

Big Cloud

Hyperscaler dependency

AI training data

AI supply chain

Fellow profile

MARGARIDA SILVA researches the political and economic power of Big Tech at the Centre for Research on Multinational Corporations (SOMO). She has published research on Amazon's marketplace dominance in Europe, Big Tech's monopolisation strategies and their pushback against the Digital Markets Act. Before joining SOMO she researched and campaigned against the political influence of Big Tech and other corporate sectors over EU policymaking at Corporate Europe Observatory.

Output

[Big Tech sets unfair terms and conditions for AI data workers globally ↗](#), Margarida Silva, March 2026, SOMO.

[The real winners of the AI Race: Amazon, Google, Microsoft, ↗](#) Margarida Silva, July 2025, SOMO.

[Why Meta's Scale AI deal must be stopped ↗](#), Margarida Silva, June 2025, SOMO.

Project purpose

- Maps the value chains of the 12 most valuable genAI startups to test whether they genuinely challenge Big Tech or are instead deeply dependent on the hyperscalers
- Positions the work as an intervention in the dominant "AI Race" / "national champions" narrative, particularly the framing of OpenAI, Anthropic, and Mistral as challengers

Key findings

- Of the 12 most valuable genAI start-ups analysed, 11 use Nvidia chips, 10 depend on Amazon, Microsoft, or Google for infrastructure, and 9 commercialise via their platforms
- Big Tech companies have leveraged their dominance in cloud to become the platform for connecting genAI start-ups with customers; Amazon Bedrock, Azure, and Google Vertex provide cloud customers with access to genAI models and tools for fine-tuning and deploying them.
- Without breaking the dependency between AI start-ups and Big Tech, public interventions in the “AI race” are likely to further entrench the power of dominant tech companies.
- The Meta-Scale AI deal represents Meta’s strategic entry into AI infrastructure, reminiscent of its acquisition of Instagram and WhatsApp; tested whether antitrust authorities would engage with labour concerns and found they would not, as these considerations fall outside antitrust authorities’ framework.
- The five largest hyperscalers use at least 30 intermediary companies for data work, several accused of paying below minimum wage, unfair dismissals, preventing collective organising, and providing no social protections; Big Tech is not just a customer but also invests in data work vendors and integrates them into its cloud marketplaces.

Key learnings

- Research relied heavily on model cards; information that was available in earlier model cards (e.g. ChatGPT, two years prior) was no longer there, indicating a reduction in transparency over time.
- Antitrust authorities’ frameworks have limitations when it comes to labour considerations.
- Future directions: creating visuals of the dependencies for civil society mobilisation, identifying policy tools such as due diligence obligations for data workers, and building strategic alliances with centre-left parties (on environmental/labour contradictions), centre-right parties (on sovereignty concerns), and business groups.

Market Power and Dependencies in the Market for Foundation Models

by Nicolas Eschenbaum



Project 7 · Individual Project

Focus

Foundation models markets

AI market definitions

Vertical stratification

Fellow profile

Nicolas Eschenbaum is a researcher and advisor for competition economics and digital markets. He studies how AI and digitisation affects markets and how digital technologies can be regulated in order to promote competition and innovation.

Output

The research has not been published and is not publicly available. Below is based on grant reporting only.

Project purpose

- Conducts an empirical investigation of switching costs and market concentration in the foundation models market.
- Motivated by the gap that through 2023-2024, competition authorities (CMA, EC, DOJ) launched investigations into AI markets, particularly examining partnerships like Microsoft-OpenAI and Amazon-Anthropic, but these discussions lacked empirical evidence on switching costs, a critical determinant of market power.
- Existing analyses relied on case studies or aggregate market data; no research captured granular technical barriers (prompt engineering debt, evaluation harness reconstruction) or contractual constraints (volume commitments, HIPAA compliance) from organisations with production deployments.

Key findings

- Surveyed 19 organisations spanning a range of firm sizes and verticals (healthcare, legal, coding, customer support) to measure switching costs
- Switching costs range from \$450 to \$2,000 for simple cases, up to \$5,000 to \$70,000 for complex pipelines
- Identifies vertical stratification as a core methodological finding: healthcare shows monopoly characteristics, coding shows oligopoly, and content generation shows moderate concentration
- Argues that competition authorities need use-case-specific market definitions rather than treating foundation models as a single horizontal market

Key learnings

- Small samples (n=19) can still generate policy impact because competition authorities value exploratory findings over multi-year large-n studies
- AI markets require use-case-specific analysis rather than assumed horizontal homogeneity

Project 8

The political economy of military general-purpose AI (GPAI) systems by Yung Au and Srujana Katta



Project 8 · Team of 2

Focus

Military general purpose AI

AI supply chains

AI training data

Fellow profile

YUNG AU is an OTF Senior Fellow at the School of Geography, University of Oxford and a Board Member of the Surveillance in the Majority World Network. Her work examines the supply chain of surveillance and policing technologies.

SRUJANA KATTA is a PhD researcher at the Oxford Internet Institute, University of Oxford. Her research interests include digital platforms, AI, and collective action in the technology sector.

Output

Note: the research has not been published and was not available yet. Below is based on grant reporting and workshop notes.

Project purpose

- Investigates the political economy of military general-purpose AI systems, applying a framework that foregrounds material conditions, power relations, and global supply chains rather than treating military AI as a purely technical or ethical problem
- Built a cross-border research network spanning Brazil, Jordan, Palestine, Nigeria, India, Hong Kong, the UK, and the US, triangulating corporate disclosures, procurement records, and investigative journalism to map deliberately opaque military AI supply chains

Key findings

- Key insight: the profit models of large-scale AI systems are not straightforward
- Militaries represent a special type of client for GPAI providers, presenting opportunities for lucrative contracts for large-scale infrastructures and affording the capture of a mandatory audience
- One case study of data centres in contested territories exemplifies the importance of looking at the materiality of tech supply chains that enable GPAI, and how physical infrastructure is used to occupy territory and “terraform the landscape”; relatedly, the military choice of where to base data centre and R&D infrastructure explicitly cites “lower ethical standards” and “less stringent regulation”
- Another case study points to the surplus value extracted from data gathered from conflict scenarios, occupied territory, and other war crimes, which is then used to feed back into AI systems
- Intended impact: input into strategic litigation efforts of civil society and union actors seeking to hold military and tech companies accountable

Key learnings

- The opacity of military AI supply chains and development was a central methodological challenge. The team triangulated corporate disclosures, procurement records, open-source investigations, legal filings, and investigative journalism. They deliberately avoided aiming for comprehensiveness, focusing instead on key categories, relationships, and illustrative cases that reveal broader political-economic dynamics.
- Working with contributors from regions affected by military AI violence was emotionally demanding and required trust-building and flexibility on timelines

ABOUT US

The European AI & Society Fund gives money and support to organisations around Europe working to ensure AI is safe and fair for everyone. Since 2020, we've pooled funding from 20 foundations to award €13.6 million to organisations across 27 countries. Together, we're shaping AI that serves people, society and planet.