

Report

AI, digital markets and competition: a research and policy agenda

Prof Diane Coyle



Author

Diane Coyle, Bennett School of Public Policy, University of Cambridge. cc700@cam.ac.uk

Suggested citation

Coyle, D. (2026) AI, digital markets and competition: a research and policy agenda, Bennett School of Public Policy, Cambridge. DOI: <https://doi.org/10.17863/CAM.136241>

Acknowledgements

This report captures the discussions in a workshop hosted in Cambridge by the Bennett School of Public Policy in July 2026, generously supported by Brian Kahin. Academics and policymakers participated, representing several disciplines: economics, law, political science, computer science and management studies. The workshop was held under the Chatham House rule. All participants have had the opportunity to contribute to this document. Those who wished to be named are listed below; we warmly thank the other participants who did not so wish.

Giacomo Calzolari, European University Institute

Diane Coyle, Bennett School of Public Policy, University of Cambridge

Hayane Dahmen, BSPP, University of Cambridge

Annabelle Gawer, University of Surrey

Oliver Hauser, Bennett School of Public Policy, University of Cambridge

Michael Kummer, Bennett School of Public Policy, University of Cambridge

Philip Marsden, College of Europe

Derek McAuley, University of Nottingham

Christine Sheldon, Bennett School of Public Policy, University of Cambridge

Jatinder Singh, University of Cambridge

Greg Taylor, OII, University of Oxford

Aleksei Turobov, Bennett School of Public Policy, University of Cambridge

Mike Walker, Durham University Business School

Andrea Renda, Centre for European Policy Studies

This report was written by humans, not by AI.

Contents

Preface	1
1. Business models and market power	2
2. Implications for policy and enforcement	7
3. (Geo)political economy	9
4. Conclusion	12
References	13

Preface

It is seven years since a series of landmark reports on competition in digital markets was published: the Furman Review *Unlocking Digital Competition* in the United Kingdom (UK); the EU's *Competition Policy for the Digital Era*; and in the United States (US) the Stigler Center's *Committee on Digital Platforms*. In each of those jurisdictions the expert analysis paved the way for legislation or judicial action. The UK passed (in 2024) the Digital Markets, Competition and Consumers Act (DMCCA) 2024. In the European Union (EU) the Digital Markets Act (DMA) was in place by 2022. During the Biden Administration Lina Khan as Head of the Federal Trade Commission and Jonathan Kanter at the Department of Justice took a more active approach to enforcement.

Much has changed since 2019. Generative Artificial Intelligence (GenAI) has dramatically altered the technological context in several ways, including new entrants (Anthropic and OpenAI), but also in terms of capabilities that are likely to reshape market structures and business models. The economics of generative AI differ from platform economics: frontier models require vast upfront capital spending, involve large economies of scale and, while not characterised by the same network effects, operate in complex ecosystems. The technology is progressing extremely rapidly, with agentic systems widely seen as the next application likely to be adopted across the economy.

The political and geopolitical context has also changed significantly. Chinese start-ups have come close to matching the technical capabilities of US frontier models, but with a commitment to smaller open weight models that can more easily be adopted. The dominance of the US and China has led many other countries to consider both their need for some AI 'sovereignty' or control, and what shape their industrial policies for AI ought to take. Competition, industrial and trade policies have converged in ways that challenge the traditional separation of policy analysis and practice (e.g. Coyle 2025, Coyle & Dahmen 2026).

We therefore held this workshop to try to identify the gaps in research and policy posed by this changed context. The discussion was structured under three headings: business models and market power; competition policy in practice and enforcement; and the geopolitical context. This document sets out the policy and research questions that arose in the discussions. The agenda that emerged is certainly not comprehensive, but we believe it identifies some key challenges generative AI poses for competition and market economies.

Diane Coyle

Bennett School for Public Policy, October 2026

1. Business models and market power

Digital markets have been characterised by dominant players for over two decades now, with tech companies occupying nine of the top ten places globally by market capitalisation. The major platforms have benefited from network effects and the data feedback loop to cement their position in growing ecosystems. They also enjoy economies of scale and scope, and advantages from vertical integration and ecosystem business models. What's more, big tech companies have been found in a number of cases to have abused their dominant position in the market, meaning the competition policy challenge goes beyond the core properties of digital markets to include current and future forms of competitive conduct. They both act as a foundation for other market activity, and as gatekeepers at key nodes in the ecosystem network.

At least two major pathways for capturing market power have emerged with generative AI development. The first is the familiar market capture of large incumbents such as Google and Microsoft that have dominated the digital markets for the past decades. These digital platforms offer users access to their AI tools through existing relationships with users in their vast multiple-product offerings. Microsoft, for example, benefits from its corporate Office 365 suite used widely across organisations, and may seek to foreclose rivals before they can get into those organisations. Similarly, Google has a large ecosystem that may offer complementarities and data loop effects, and is the default access point to AI in Android and search.

The second, and more remarkable, is that two key new entrants, Anthropic and OpenAI, have rapidly ascended to become important players in the market for AI. Yet each is still trying to crystallise a viable business model. OpenAI has millions of free users it aims to upsell into paying customers. Anthropic is focusing on enterprise uses and subscription services; at this point it has now reported two successive quarters of profits. Both Anthropic and OpenAI have received large investments from the big tech companies. The landscape of US-based AI companies indeed involves multiple connections between the major players, including equity stakes, partnerships, and access to compute (<https://ai-circular-economy.com/>). There have in addition been a number of acquisitions (such as HuggingFace by Nvidia, and ScaleAI by Meta). They are joined, increasingly, by a number of Chinese competitors.

The frontier labs need to operate at daunting scale, given their focus and strategic choices. The upfront cost of research and model training, the need for data and advanced chips, the data centre and energy/water requirements make the underlying economics more like infrastructure than platform models. Some of the frictions that led to market power in platform markets remain, such as switching costs for enterprise software or network effects for application programming interfaces (APIs), as AI is a label for a range of different products and markets with differing characteristics. But the level of capital expenditure required (and hence supply-side economies of scale) is a new and significant feature, one not yet bypassed by technical

developments (such as highly capable small language models requiring less compute and energy).

Although venture capital has flooded into AI, the leading companies have created a complex web of cross-shareholdings and in-kind transactions (<https://ai-circular-economy.com/>). They are oligopolistic rivals competing fiercely with each other on some frontiers, and at the same time sometimes collaborating, in a 'frenemy' dynamic.

There are certainly some signs of competitive effects, indicating that the foundation model layer may be relatively contestable. Network effects are weaker. Chinese models have been adopted by some major US and European enterprises (such as DoorDash, Airbnb and Siemens; see for example, CNBC on 7 July 2026). Unified API gateways and marketplaces, most prominently OpenRouter, are emerging; its leaderboard shows increasing use of open weight models. So, some switching therefore seems to be occurring, although the providers may nevertheless be able to lock in customers through systems integration and tailored 'scaffolding' – that is, engineering in switching costs as a strategic choice. Price per token has been falling (although token use has been rising with successive models, such that price per task has declined relatively little, Zhu 2026). Rivalry also spurs the continuing efforts to push forward the innovation frontier, now focused on agentic systems as well as model capabilities. The desire by competing frontier labs to be seen as innovative and producing the most capable AI model has also spurred the development of potentially unsafe models in an unregulated environment (e.g. recent examples of Anthropic and OpenAI models 'escaping' safe testing environments, hacking third-party organisations).

At the same time, the big companies are increasingly vertically integrating, from data centres and chips up through the layers of the tech 'stack' to applications which may be bundled with existing software or linked through complementary services and APIs. New entrants have also seen opportunities for vertical integration, especially downstream. This is particularly true for agentic systems that offer new opportunities to integrate models, agents and payment rails. Vertical integration offers a playbook of anti-competitive behaviours, such as self-preferencing, exclusive tying, raising rivals' costs.

Generative AI and the current and evolving market structures raise multiple questions: AI markets are ecosystems, whereas competition analysis and practice are framed around markets, so the reality does not fit the current analytical boxes. What analytical tools will help competition authorities understand ecosystem-based competition, both between ecosystems and within them? This implies a more fundamental research question about the underlying economic welfare analysis: what does a good outcome look like? How should 'good' be understood and measured? Current regulatory tools cannot address the form of 'polycentric' competition and 'frenemy' behaviour apparent now.

How should the analysis incorporate the way market power operates across layers of the stack, and the moves towards vertical integration? Where is value being created and where is it being captured in algorithmic value chains (Cobbe et al. 2023)?

Control of agentic systems also seems set to become important, and likely to pave the way for new business models. The incumbents have significant advantages including feedback loops, adjacent services such as payments, insurance and dispute resolution. A particularly important potential source of lock in is the representation of the user's preferences. A personal agent may accumulate a rich mapping of preferences, history, constraints and objectives that is costly to recreate elsewhere. Portability may therefore need to extend beyond raw personal data to memory, and preference representations (a difficult task).

The specific competition challenges involved in Personal Agentic AI include:

- AI deployed by dominant firms with access to granular consumer data to raise firms' predictive accuracy gives them superior results with positive feedback to their market position
- Agentic systems operate as black boxes, so difficult to monitor and audit, and - at present - no responsibility for agents' actions attached to the provider
- Agentic systems operate continuously and adaptively, making them hard to monitor and audit; it may be impossible to identify their actions, roles or even existence
- They provide hyper-personalisation and proactive support but potential loss of consumer agency
- Consumers face the risk of manipulation serving the interests of third parties (potentially embedding sponsored content into AI-mediated experiences)

Hyper-personalisation could mean a better shopping experience as consumers lose some agency in day-to-day consumption decisions, but benefit from having an all-seeing shopping assistant. Agents will potentially solve problems like search frictions, choice overload, dark patterns, and switching costs while at the same time creating the kind of problems discussed here. Its spread will also change the way tech platforms operate. Without human attention to compete for, the entire advertising revenue business model is likely to change in fundamental ways.

How will welfare in personalised AI-mediated transactions be measured? When users' granular data is fed to agents and companies for greater personalisation, it could capture better value or be manipulated by companies for their gain – how would antitrust disputes measure consumer welfare where each service offering can be uniquely personalised? Might survey based methods (e.g. Brynjolfsson et al. 2025) be useful ways to generate evidence on welfare gains?

Implicit collusion by AI tools has been shown to lead to prices above competitive levels. How will collusion be avoided when algorithms implicitly learn and coordinate prices without any directive to do so by companies using the algorithms? Algorithms could also collude in ways

that are less visible and harder to measure than prices. For instance, by coordinating personalised recommendations, by synchronising stock outs, or by coordinating the targeting of ads.

Conversely, could agents on the demand side alter the balance of power? What market infrastructure would enable this – can principles of market design shape good outcomes in agentic markets? Or facilitate competition between integrated agentic systems?

With agentic transactions, rules governing information flows, communication, matching, negotiation and execution become part of the market institution itself. Competition policy may need to engage increasingly with market design and market microstructure rather than treating the environment as given.

What is the analytical framework to understand consumer outcomes when consumers are not making the immediate transaction choices?

Is data portability for users technically possible and behaviourally likely; could it pre-empt the formation of walled agentic systems (in contrast to how mobile operating systems and app stores formed)? Can proprietary data enable some firms to maintain an advantage over AI providers? Understanding when data remains a complementary asset and when their value can be absorbed into a general model is important for predicting vertical disintermediation or entry barriers.

How is innovation to be integrated into competition analysis, and does the form of the innovation make a difference? (This is not an AI-specific question, but a broader issue for competition policy.)

If AI markets are more like infrastructure markets, what kind of regulatory model would be effective? Economic regulation is hard enough with traditional infrastructure - because of information asymmetries and the risk of capture - and seems much harder with a fast-moving and general purpose technology. How would solutions such as access mandates – and the associated access pricing – operate?

And to the extent that AI is like public infrastructure, which part of the 'AI stack' (if any) would need to be operated as a public good by the state, or as a privately operating market? What would a regulator's remit and role be in creating an effective competitive environment and ensuring access to the broader public to this infrastructure? What is the role for standard-setting and interoperability? The economics of public infrastructure further intersect with issues on sovereignty (see also: geopolitics section).

Turning to other markets that will be affected by AI, one concern is that generative AI systems will disintermediate the services offered by other companies such as SaaS (Software as a

Service) providers or professional services such as law and advertising (OECD 2025). The risk varies by sector, but is non-zero, as seen by the share price performance of some of the companies potentially affected. If true, this may pose a particular economic risk to one of the UK's traditionally strongest labour market segments, the professional service sector. How big is this risk and how should it be evaluated in terms of policy, given the inevitability of structural change?

Does this risk make self-preferencing an even bigger issue than it is already? Does the existing legislation suffice for the new challenges? Should regulators force AI companies to opt for either supplying such markets or operating themselves in such markets, in the kind of split enforced in markets such as financial advice?

What will the impact of AI be in other markets? Some upstream markets providing inputs to AI are far more concentrated than downstream markets, due to very large economies of scale and their specialised knowledge/research & development. Globally there are no more than five advanced chip fabricators, with only TSMC providing the most advanced, and about 20 chip designers. The AI supply chain is extremely complex and global (Liu & Turobov, 2026 Bennett School working paper forthcoming), and the upstream concentration speaks to the issues of sovereignty discussed below.

In downstream markets, current evidence suggests that big and productive firms are more likely to have adopted generative AI. Will this continue, intensifying the existing trend to more concentrated markets in many economies? Or – given the strand of literature that theorises that generative AI will enable much smaller growth companies in future – might the technology offer SMEs new opportunities? And what of the potential for vertical absorption – where an upstream provider leverages the insider knowledge it gained by serving downstream customers to then compete against them?

2. Implications for policy and enforcement

The discussion noted that developments in AI markets have broad implications for other policy regimes, including intellectual property, liability and contract law, innovation policy, and industrial and trade policies (Coyle & Dahmen 2026). These were beyond the scope of the discussion.

On competition policy, there was broad agreement that traditional competition enforcement faces serious challenges of speed, information and technical capacity in AI markets; some participants considered that existing laws and regulations are struggling to enforce competition in AI markets. One view was that the political process has been ‘captured’ by incumbent tech companies but in any case, many agreed that competition law has been struggling in this context, partly due to slow action on enforcement, and partly due to the fact that it took too long for authorities to update their analysis after the economic literature had been pointing out the features of digital markets. Together, these mean that authorities cannot act with the necessary speed, nor do they have the requisite toolkit, meaning they face some challenges regarding these rapidly changing markets.

Current legislation predates generative AI. There has been little enforcement to date in AI markets. Indeed, there has been relatively little enforcement in existing digital markets at all since the new legislation, although this is beginning to change in some jurisdictions in, for example, cloud markets and app stores. There are no specific legal powers in all jurisdictions to intervene in practices common among the tech companies that clearly are potentially anti-competitive in their effects, such as partnerships, investment agreements, barter agreements and acqui-hires. Given that competition authorities are widely thought to have been slow to respond to past phenomena in digital markets, and that even after 2019 new legislation and court action took a considerable time, there are important practical challenges.

First, when is the right time to intervene in nascent markets (such as agentic)? Should competition authorities (or legislators) ‘move fast and break things’, on the precautionary principle, if we think the lesson of history in digital markets is that entrenched positions are difficult to reverse? Can more forward-looking analytical tools be developed? In fact, can and should regulators use the same frontier AI models against the companies developing them, to strengthen their enforcement capabilities? For example, with scenario analysis, authorities could identify a small number of plausible alternative market structures and bottlenecks, with the purpose of avoiding waiting until market power has already become entrenched before analysing possible interventions. Does this dependence on such AI tools in markets, which they are in charge of regulating, make them stronger and more able to keep up, or more fragile and dependent (e.g. “bite the hand that feeds you”)?

Given that there still seems to be competition for the market, do older lessons about interventions into competition for markets carry over?

How can competition authorities monitor or audit digital and AI markets, even for existing phenomena such as algorithmic price collusion? Can counterfactual explanations help (Wachter et al. 2018)? What resources, capabilities and technical tools would be needed for swifter and more dynamic auditing of market phenomena and intervention? Can enforcers ever hope to match the companies?

What process changes are needed? What political mechanisms should be adapted for/to AI and how? How can we navigate such adaptation in mechanisms and processes? What role and form of enforcement can be feasible to achieve? Competition cases take years whereas companies can change models almost overnight, making redundant investigation of their prior conduct.

If technical tools prove better at screening for anticompetitive behaviour, what are the right thresholds or criteria to keep caseloads manageable? How would collusion between algorithms be defined when no human or company-led intention to collude can be detected in the process of installing the algorithms? What information retention and record keeping should be required of companies?

Information asymmetries between companies and enforcers are extreme, and companies may be able to keep key information from competition authorities. This can include proprietary training data, algorithms, and personalisations that steer the model towards the company's own benefit at a consumer's loss. Can an auditing regime counter this – can access to independent researchers be mandated? Which part of the AI stack would need to be made available for accessing? Can civil society organisations play a role in monitoring?

Are behavioural remedies ever sufficient or should the authorities be looking to divestment or other structural remedies? What is the role of pre-assigning liability for consumer harms?

Why have competition authorities so rarely used the tool of interim measures in digital markets, to help overcome some of the procedural problems, namely interim measures. The recent case regarding Meta and its WhatsApp AI chatbot may signal a change.

Is competition policy even an adequate tool – should AI markets be covered by a more active regulatory regime, as in telecoms or banking? Or is a flexible approach (such as the UK's DMCC Act) preferable as it gives flexibility to respond to a changing technology and markets.

Looking at the impact of agentic systems, there is a large agenda to understand their potential impact on market structures. Among the priorities would be traceability of their actions, liability, and identity systems for agents. Will agentic systems be able to accurately capture and map consumer preferences into actions? How would this be monitored and enforced?

3. (Geo)political economy

The issues under the heading of political economy fall into two categories. The first is the role of competition authorities in the context of the technological transformation that is reshaping market structures, and the perceived failure of competition or antitrust law to tackle digital market power to date. The second is the geopolitical dimension, given the framing of frontier AI as a 'race' between the US and China, and in which the EU, as well as the UK and other non-great powers, are seeking to catch up – or at least keep up with in some way.

On the first, an institutional question is the extent to which competition authorities can or should become involved in areas not previously in their remit. Given the convergence of (at least) competition, industrial and trade policies, can those remits be clearly specified at all – for example, what if a foreign acquisition would affect domestic capabilities even if it did not affect competitive dynamics?

Is the focus in many jurisdictions on consumer welfare an appropriate analytical frame? What is an appropriate normative framing for policy? To what extent should potentially different levels (and potentially competing) welfare criteria be prioritised – for instance, AI may help consumers individually (e.g. better advice, better service) but may have societal downsides (e.g. systemic risks or fragility)?

Should authorities stay narrow in their focus on the ex-post application of the law, and leave competition policy including more forward-looking tools to others? Or alternatively, are they the right institutions with the expertise to expand to a wider remit. Is there a trade-off between authorities' independence, suggesting they should not intervene in policy choices with a normative or political dimension, and their relevance?

Given the widely shared perception of the inadequacy of competition law and policy to date, certainly in the context of digital markets and AI, and possibly beyond, what might an active regulatory regime for the AI sector(s) involve? Does the existing set of legislation and regulation such as the DMA and DMCCA offer tools to be more proactive in AI markets?

The societal need is increasingly being described in terms of digital public infrastructure, which integrates the multiple levels of the AI stack itself along with associated elements such as identity (human and agentic) and potentially also services such as payment rails or public services. What is infrastructural, what is potentially competitive, and how can value creation and capture be conceptualised? How can commercial logics and the need for shared assets be combined, and in what economic welfare framework? How might generative AI providers interface with digital public infrastructure? What is the need for pro-active measures to increase AI-literacy?

One contentious area with no settled analytical or policy approach is the role of data – itself a key layer in digital public infrastructure. The frontier US AI labs accessed significant amounts of copyright as well as open data to train their models in the first instance. The Chinese labs have developed models using the same data through distillation techniques. Legal cases brought by copyright holders are ongoing, while the AI companies are also negotiating commercial access to material with publishers. Other examples of disputes over copyright include companies (such as GM, John Deere) claiming ownership over data produced by their customers and aggregated for licensing back to those customers. The copyright questions are by no means resolved. A fundamental challenge is developing an approach to estimating the value of the data, given the positive (information spillovers) and negative (privacy spillovers) externalities involved.

A further long-term challenge is how to continue to incentivise the production of the kind of data that made AI possible in the first place? Creators and innovators are now at risk of losing their comparative advantage and income streams, as their data is being fed to build models that then help to power similar capabilities and outputs in future models. From a technical perspective, it may be necessary for new data to be generated by humans to avoid ‘model collapse’ (i.e. degradation of model performance) resulting from algorithms being trained (too extensively) on their own output.

Moving on to the second aspect of political economy, what active industrial and trade policy and competition policy mix should policymakers adopt to ensure AI markets that are both competitive and sufficiently ‘sovereign’, that are able to be shaped in ways that serve the domestic market? Is there any reality to the idea of a completely ‘sovereign’ AI stack in any country, and how much can and should this narrative be carried into a feasible trade policy? If not all aspects of the AI stack can be internally created, which seems inevitable, what will be the minimum expectations and requirements to be sufficiently sovereign as a nation? While some sensitive activities may require UK entities and infrastructure, in many contexts the answer may be developing a range of mutual partnerships with like-minded countries, with a view to minimising specific strategic vulnerabilities.

The metaphor of a race between the US and China is somewhat misleading. While the big US labs are racing each other to stay at the technological frontier, China has concentrated on some catch up and at the same time on broader deployment of AI applications. The recent [intervention by Xi Jinping](#) was a clear signal of China’s intention to use its open-weight approach to expand its AI models in overseas markets.

The dominance of the frontier by the US and China is widely seen as a reason for the EU, UK and other powers to develop their own AI strategies. It should be remembered that for both the frontier countries there has been significant state intervention to support their domestic companies. What constitutes an effective combination of industrial and competition policies for these other markets? How should supply chain dependencies be incorporated or navigated?

Recognising that AI has the potential to become the transformational technology of the 21st century, what can we learn from the US and Chinese regulatory approaches? Can we locate best practices in these and potentially other markets? For example, Korea has provided its entire legal corpus and government publications in .md format for AI to be able to seamlessly access the information and better understand regulation. Are other countries leading with similar useful practice for both regulating AI, but also for encouraging adoption?

Where countries want to build a reputation in their national AI capabilities, should they focus on one or multiple layers of the AI stack? What is the role of human capital, both in terms of existing talent (e.g. skilled AI engineers; AI developers for downstream applications) and formation of talent (e.g. universities, research labs, frontier companies themselves)? How does human capital influence a country's geopolitical economy and AI strategy when AI talent is globally in high demand and recruited internationally?

4. Conclusion

The technology of generative AI is advancing quickly, markets up and down the AI stack are fluid, the supply chains are complex and global, many transactions are taking place, and the political and policy environment is turbulent. So, it is not surprising that this note has set out many questions. The common theme is that AI poses a significant challenge to competition analysis, policy and enforcement (and indeed to other policy domains as well). The sense among participants in the workshop was that analytical tools and enforcement actions were slow to catch up with earlier developments in digital markets. Although the EU and UK have since implemented new legislation, the dynamics of generative AI will pose new challenges for researchers and policymakers.

References

- Brynjolfsson, E., Collis, A., Diewert, W.E., Eggers, F. and Fox, K.J. (2025) 'GDP-B: accounting for the value of new and free goods', *American Economic Journal: Macroeconomics*, 17(4), pp. 312-344. doi: 10.1257/mac.20210319.
- CNBC (2026) *Chinese AI models are gaining ground with U.S. companies as OpenAI, Anthropic costs surge*. 7 July. Available at: <https://www.cnbc.com/2026/07/07/chinese-ai-models-costs-us-openai-anthropic.html> (Accessed: 7 August 2026).
- Calvano, E., Calzolari, G., Denicolò, V. and Pastorello, S. (2023) '[Algorithmic collusion: genuine or spurious?](#)', *International Journal of Industrial Organization*, Elsevier, vol. 90(C).
- Calzolari, G. and Hanspach, P. (2025) 'Pricing algorithms out of the box: a study of the repricing industry', *Journal of Competition Law & Economics*, Volume 21 Issue 2, pp. 163-187. <https://doi.org/10.1093/joclec/nhaf003>
- Cobbe, J., Veale, M. and Singh, J. (2023) 'Understanding accountability in algorithmic supply chains', in [FAccT '23: Proceedings of the 2023 ACM Conference on Fairness, Accountability, and Transparency](#). New York: Association for Computing Machinery, pp. 1186-1197. [doi:10.1145/3593013.3594073](https://doi.org/10.1145/3593013.3594073).
- Coyle, D. (2024) 'Everything everywhere all at once: competition policy and industrial policy choices in an era of structural change', *Oxford Review of Economic Policy*, 40(4), pp. 718-728. [doi:10.1093/oxrep/grae040](https://doi.org/10.1093/oxrep/grae040).
- Coyle, D., and Dahmen, H. (2026). When worlds collide: The asymmetric convergence of competition, trade, and industrial policy. *Oxford Review of Economic Policy*. (Forthcoming)
- Liang, K., Hu, H., Zhao, X., Song, D., Griffiths, T.L. and Fisac, J.F. (2025) 'Machine bullshit: characterizing the emergent disregard for truth in large language models', *arXiv preprint*, arXiv:2507.07484.
- Liu, S. and Turobov, A. (2026) *The invisible supply chain*. [Forthcoming, to be published in November 2026]. Bennett School of Public Policy.
- OECD (2025) *Artificial intelligence and competitive dynamics in downstream markets*. Paris: OECD Publishing. Available at: https://www.oecd.org/content/dam/oecd/en/publications/reports/2025/11/artificial-intelligence-and-competitive-dynamics-in-downstream-markets_c6e81d0e/ccf0624a-en.pdf (Accessed: 7 August 2026).

Peters, U. and Chin-Yee, B. (2025) 'Generalization bias in large language model summarization of scientific research', *Royal Society Open Science*, 12(4), article 241776.

Salecha, A., Ireland, M.E., Subrahmanya, S., Sedoc, J., Ungar, L.H. and Eichstaedt, J.C. (2024) 'Large language models display human-like social desirability biases in Big Five personality surveys', *PNAS Nexus*, 3(12), article pp. 533.

Wachter, S., Mittelstadt, B. and Russell, C. (2018) 'Counterfactual explanations without opening the black box: automated decisions and the GDPR', *Harvard Journal of Law & Technology*, 31(2), pp. 841-887. Available at: <https://arxiv.org/abs/1711.00399> (Accessed: 1 August 2026).



UNIVERSITY OF
CAMBRIDGE



Bennett School
of Public Policy

Bennett School for Public Policy
Mond Building
New Museums Site
Free School Lane
Cambridge CB3 9DT

bennettschool.cam.ac.uk
enquiries@bennettschool.cam.ac.uk

