Autorité de la concurrence

"The simple macroeconomics of transformative AI" by Daron Acemoglu

Comments by Benoît Cœuré, Autorité de la concurrence[1]

Economic Policy Panel, Brussels, 4 April 2024

Thank you for inviting me to discuss this important paper.

Let me first say that our discussion couldn't be timelier. Artificial intelligence (AI) is everywhere in the public debate. It is fair to say that it raises expectations which are disproportionate to its current economic significance. Calibrating – even roughly – its macro implications will help budget the time that policymakers should spend on the issue.

Daron Acemoglu's paper does a very good job at providing us with a consistent macro framework to describe the channels through which AI impacts the economy. It is an important addition to a burgeoning field of studies.[2] It is reasonable and conservative: reasonable in that it accounts for equilibrium effects on the labour market, and conservative in viewing AI as incremental technological change rather than a quantum leap.

The description of the labour market impact of AI is compelling, with important conclusions such as the impact of AI being greater on less-educated females, while previous innovations have first impacted less-educated males. I am tempted to cut through the intricacy of the model and view this conclusion as a consequence of AI primarily having an impact on the provision of services (heavily skewed towards women), while previous innovations have impacted manufacturing jobs.

The conservative approach produces numbers – a 0.7% impact on total factor productivity in the US economy, a slightly higher than 1% impact on US GDP – which stand at the low end of existing estimates, particularly coming from the private sector.[3] This is enough to allow for upside surprises in the future, and to spur a good discussion now.

Tonight I would like to start from Daron's framework, and wonder what could go wrong and yield an even lower impact.

Modest as they are, the productivity gains identified in the paper are conditional on two assumptions: the AI industry staying the innovation course, and AI-driven innovation being adopted and exploited across the broader economy. The bad news is that concerns have already been identified and are being actively scrutinised by competition authorities.

So, where are the choking points?

---

[1] All views expressed here are mine only.
[2] See e.g. Aghion and Bouverot (2024), Besson et al. (2024), Bryjnolfsson and Unger (2023) and Council of Economic Advisers (2024).
[3] As an example, Briggs and Kodnani (2023) foresee a 1.5% increase in yearly labour productivity *growth* over a 10-year period, and a cumulated impact on global GDP of 7% thanks to generative AI.

Let me address first risks to competition along the AI value chain, then in the broader economy. To conclude, I will say a few words about the impact of AI on the economics of data.

Risks to competition along the generative AI value chain

Many of these concerns relate to the fact that the inputs needed upstream to design, train and fine-tune Large Language Models (LLMs) are costly and already partly in the hands of globally active big tech companies. When one compares AI with earlier technological advances, one should bear in mind that this is a rare instance where a potentially disruptive innovation is controlled at an early stage by incumbents.

These upstream inputs are, *inter alia*: data, computational power ("compute" in industry jargon), the supply of cloud services, highly specialised skills, and of course, money.

They entail increased returns to scale due to a powerful mix of high entry costs (e.g. capacity investment in data storage and GPUs),[4] pecuniary externalities (e.g. due to the breadth and depth of AI and cloud marketplaces), network effects, and learning-by-doing à la Arrow (1962) through training on ever expanding datasets.[5]

Controlling these inputs may allow big tech companies to cement their market power, foreclose the market for LLMs, and extract value from downstream users, e.g. by bundling AI and other services.[6] In antitrust jargon, there are risks of both exclusionary and exploitative abuse.

This is not a theoretical discussion. The market for cloud services already provides a cautionary tale of what could go wrong.[7]

There are other issues associated with market power, not least of which risks to the democratic process. This includes risks of manipulating elections, steering public opinion and capturing the regulatory process. Protecting democracy should very much be on antitrust enforcers' minds.

Competition authorities all agree that strict antitrust oversight will be needed. The European Commission has collected feedback from industry participants and is investigating financial partnerships between big tech companies and AI startups. At the French Competition Authority, we have launched an *ex-officio* inquiry into the upstream generative AI value chain, to be published end-June.

Another area of concern is a lack of transparency in the collection and use of data used to train and fine-tune LLMs. Being publicly available does not mean being free of rights. The French Competition Authority recently fined Google €250 million for (*inter alia*) not informing French publishers and press agencies of their content being used to train Google's conversational AI service, Bard (now called Gemini), and for not providing them with the effective opt-out foreseen by French and European law.

---

[4] The cost to train Google's Gemini has been publicly estimated to be more than $600 million.
[5] See Crémer et al. (2019) and Furman (2019) for surveys of increasing returns to scale in the digital sector.
[6] See also Gans (2024) and Korinek and Vipra (2024).
[7] See Autorité de la concurrence (2023).

Turning to the downstream generative AI value chain, where LLMs are used to provide AI services to individuals and companies, there may be room to be (cautiously) more optimistic for future competition.

The variable cost of using an LLM (aka "inference") is much smaller than the fixed cost of training it. There will be a variety of use cases. Different types of data will be used to fine-tune LLMs to specific needs: manufacturing data, health data, academic data, and so on. Part of this data will be proprietary and will not entail the same learning-by-doing effects as data used to train general-purpose AI. Sectoral regulatory requirements will also limit the scope for increasing returns to scale, and sectoral AI markets will be thinner.

One should, however, be wary of the vertical effects I just mentioned, as well as of conglomerate effects – acquired market power in AI allowing big techs to expand in other markets, and vice versa.

Protecting competition in the AI sector can spur faster adoption by individuals and companies, and thereby a faster diffusion of productivity gains. Adoption entails transaction costs within companies (employee training, process innovation, and so on). Powerful AI providers may be tempted to design strategies to lock-in clients and make switching between AI providers and/or multi-homing costlier, which would slow down adoption. Using cloud services, which are inherently scalable, may relieve companies of the need for upfront IT investment, but may expose them to potentially harmful practices by cloud hyperscalers.[8]

Risks to the broader competitive process

Reaping the macroeconomic benefits of AI identified in Daron's paper also assumes that AI will not hamper (and ideally will spur) the competitive process in the broader economy. It is our duty as competition enforcers to ensure that efficiency gains from AI are not only maximised, but also fairly spread across the economy. Beyond the AI value chain, we need to keep an eye on AI use cases and how they will change industry structures.

There are many reasons to be optimistic. On the supply side, the prospect of plugging AI into production processes should spur competition on the merits between companies producing similar goods and services. On the demand side, although it is difficult to be specific without knowing more about future use cases, one can also envisage a future where AI adoption will spur competition on diversity and quality.

AI could also benefit consumers by allowing them to harvest and select information which is more relevant to their needs and consumption patterns, by making it easier for firms to deviate from collusive equilibria,[9] and by allowing competition authorities to better detect cartels, all of which would be pro-competitive. We know, however, that algorithms can increase risks of tacit collusion.[10]

---

[8] Autorité de la concurrence (2023), op. cit.
[9] See Miklos-Thal and Tucker (2019).
[10] See the joint study by the Autorité de la concurrence and the Bundeskartellamt (2019).

Generative AI and the economics of data

One "known unknown", on which I would love to see more theoretical and empirical work, is the impact of generative AI on the economics of data. As emphasised by De Cornière and Taylor (2022), whether the use of data by companies makes markets more or less competitive is a difficult question because firms use data in many different ways (targeted advertising, price discrimination, product improvement, etc.), and because hoarding data entails the increasing returns to scale I already described.

Data is generally assumed by economists to be non-rival. The rise of generative AI may create a new universe of protected content: still non-rival, but excludable. This is either because firms fine-tune LLMs with proprietary data, or because they enter into exclusive deals with LLM manufacturers, as some press publishers have already done.

What are the policy consequences?

Issues related to data interoperability or functional equivalence will arise in AI as in other digital industries. There do not seem to be major obstacles today for firms to multi-home (i.e. to use different AI services at the same time, e.g. Copilot together with Gemini), but we will need to keep an eye on this in the future, and act as needed using regulation.

A key question will be the balancing act between incumbents' wish to protect data and data-related skills, the need to maintain access to allow for new entrants, and privacy and copyright requirements. A related question will be the impact of steeply increasing demand for data on the price of data, and the risk of creating another barrier to entry. Addressing these questions will require better understanding of the role of data quantity vs quality in training and fine-tuning LLMs, and of emerging techniques to reduce the data footprint of an LLM, such as using synthetic data or self-training.

I suspect that the economic consequences will in any case be massive.

* * *

Competition agencies, much like anyone else, have no crystal ball to predict the future. But what we can do is set the scene to make sure that AI is a force for growth and innovation, where smaller players have a chance to succeed and where companies and consumers have access to a variety of models. This will be our modest contribution to the modest gains identified in this fascinating paper.

References

Acemoglu, D. (2024). The simple macroeconomics of transformative AI. Presented at the 79th Economic Policy Panel Meeting, Brussels, 2-5 April.

Aghion, P., and A. Bouverot (2024). IA: Notre ambition pour la France. Rapport de la Commission de l'intelligence artificielle, March.

Arrow, K. (1962). The economic implications of learning by doing. Review of Economic Studies, 29 (3), 155-173.

Autorité de la concurrence and Bundeskartellamt (2019). Algorithms and competition, November.

Autorité de la concurrence (2023). Opinion 23-A-08 on Competition in the Cloud Sector, 29 June.

Besson, L., A. Dozias, C. Faivre, C. Gallezot, J. Gouy-Waz, and B. Vidalenc (2024). Les enjeux économiques de l'intelligence artificielle. Trésor-Eco, 341, April.

Briggs, J., and D. Kodnani (2023). The potentially large effects of artificial intelligence on economic growth. Goldman Sachs Global Economic Analyst, 26 March.

Brynjolfsson, E., and G. Unger (2023). The macroeconomics of artificial intelligence. Finance & Development, December, 20-25.

Council of Economic Advisers (2024). An Economic Framework for Understanding Artificial Intelligence. In Economic report of the President 2024, Chapter 7, 243-290.

Crémer, J., Y.-A. de Montjoye, and H. Schweitzer (2019). Competition policy for the digital era. Final report for the European Commission, Directorate-General for Competition.

De Cornière, A., and G. Taylor (2022). Data and competition: A simple framework with applications to mergers and market structure. CEPR Discussion Paper, DP14446, 8 March.

Furman, J. (2019). Unlocking digital competition. Report of the Digital Competition Expert Panel, March.

Gans, J. (2024). Market power in AI. NBER Working Paper, 32270, March.

Korinek, A., and J. Vipra (2024). Market concentration of Foundation Models: the Invisible hand of ChatGPT. Presented at the 79th Economic Policy Panel Meeting, Brussels, 2-5 April.

Korinek, A., and J. Stiglitz (2019). Artificial Intelligence and its implications for income distribution and unemployment. In Agrawal, A., J. Gans, and A. Goldfarb, The Economics of Artificial Intelligence: An Agenda, University of Chicago Press.

Lagarde, C. (2023). A Kantian shift for the Capital Markets Union. Speech at the European Banking Congress, Frankfurt-am-Main, 17 November.

Miklos-Thal, J., and C. Tucker (2019). Collusion by algorithm: Does better demand prediction facilitate coordination between sellers? Management Science, 65 (4), 1552-1561.

Varian, H. (2019). Artificial Intelligence, Economics, and Industrial Organization. In Agrawal, A., J. Gans, and A. Goldfarb, The Economics of Artificial Intelligence: An Agenda, University of Chicago Press.