Book review: The Microeconomics of Artificial Intelligence

Maximilian Kasy, February 2026

How should economists think about the impact of AI? One option is to understand technology in general, and AI in particular, as a *production function shifter*. As such, AI might differentially impact the marginal product of different factors of production, and in particular of different types of workers, which leads to concerns about skill biased technical change, polarization, automation, etc. This perspective has been very fruitful for structuring and guiding empirical research in labor economics, but it is not specific to AI.

Alternatively, using a more specific approach, we might model AI as a source of *information*, that is, of signals about an unknown state of the world. The new textbook *Microeconomics of AI* by Joshua Gans takes this latter approach. The focus throughout this book is on human decision-makers facing uncertainty. The unknown state of the world impacts the payoff of decisions. AI can provide a signal about this state, and thus about payoffs, which allows humans to make better decisions.

This modelling approach yields many interesting corollaries, which *Microeconomics of AI* spells out carefully. Information provided by AI can act as a substitute to alternative forms of risk mitigation, such as insurance or protection, and can act as a complement to alternative sources of information (judgements) about payoffs. These models shed light on the tradeoffs between full automation and human-in-the-loop decision-making, including the attendant agency problems. Understanding AI as information can also refine our understanding of its potential for automating tasks.

The book further explores how AI is produced and priced, focusing on the central role of data. It shows how incentives to provide data shape the quality and availability of AI systems, and how feedback effects can create new challenges over time. These forces influence how AI is priced and sold, how competition unfolds, and when market power or coordination problems arise. The final chapters examine policy issues such as competition, privacy, copyright, misinformation, and algorithmic bias.

Microeconomics of AI is a textbook aimed at advanced undergraduate or graduate students in business and economics. As such, its guiding motivation is that of business: How should AI be used for commercial success, and how do processes and organizational structures need to be revised in the age of AI? The book's primary tools are those of applied microeconomic theory. This book provides an excellent introduction to the way microeconomists think about AI, by discussing a coherent set of stylized economic models that speak to a wide range of important issues.

What might be missing from this book, and where should the literature go from here? Models of *AI as information*, as in this textbook, are more specific than models of *AI as production function shifter*, but the level of abstraction of such models still remains at quite some distance from the specifics of machine learning and AI. I believe it would be very productive for economists to engage more deeply with the literature in machine learning itself. There is a lot of natural overlap between the frameworks of machine learning and those of economic theory. Both fields use conceptual tools such as optimization, probability, and decision theory, which makes productive interaction between these fields quite possible. The starting point for all of machine learning, and for AI more broadly, is the specification of decision-problems, with objectives (rewards or losses), action spaces, and sources of information. Algorithms in AI are thus always constructed as automated decision-makers; typically (and in contrast to the book's approach) without humans in the loop.

The machine learning literature has generated many insights of direct relevance for the economics of AI, and more specifically for the topics discussed in *Microeconomics of AI*. A couple of examples help to illustrate:
The literature on *scaling laws* provides detailed quantitative insights on the *production function of AI*, both empirically and theoretically, across application domains. Well-developed statistical theory on variance/bias tradeoffs and tuning relates predictive performance to the amount of data and compute available. Empirical explorations of this relation (cf. Kaplan et al. (2020), Hoffmann et al.

(2022)) have motivated the "bet on scaling" in language modelling.
As another example, the literature on statistical *learning subject to differential privacy* has proven mathematically that a variety of learning problems (in supervised learning, multi-armed bandits, etc.) can be solved with essentially no loss to performance, while maintaining differential privacy for individuals contributing data (cf. Dwork & Roth (2014), chapter 11). As a consequence, individual data-sharing can be be made incentive compatible at low cost regardless of the downstream consequences for the individuals in question. These results provide strong quantitative proof of *data externalities*: Data externalities are not just an accidental feature of some domains of AI, but are instead inherent to the very nature of machine learning as the discovery of patterns *across* units of observation.

On a separate note, there is the question of how AI objective functions are chosen and who gets to choose them. *Microeconomics of AI* is to be applauded for its emphasis on judgement, interpreted as evaluation of the payoffs associated with different outcomes, in the context of human-in-the-loop decision-making. What should, I believe, receive more attention is that these judgements also involve conflicts of interest and values between different people, around the objectives that are maximized by AI algorithms. Consider the filtering of information by machine learning-based algorithms on social media and search engines, for example. These are typically designed to maximize engagement and ultimately revenues from ad clicks. But as a society, we might rightly worry about the attendant consequences for democracy and mental health, for instance. The gap is arguably not so much a matter of judgement (by engineers at Facebook and Google), and more a matter of whose objectives the algorithms are designed to pursue.

None of these considerations take away from the impressive achievement that is *Microeconomics of AI*. Instead, these points should be taken as good news for the curious economist: much exciting work remains to be done, building on the literature that is so clearly discussed in *Microeconomics of AI*.

References

Dwork, C., & Roth, A. (2014). The algorithmic foundations of differential privacy. *Foundations and Trends in Theoretical Computer Science*, *9*(3–4), 211–407.

Hoffmann, J., Borgeaud, S., Mensch, A., Buchatskaya, E., Cai, T., Rutherford, E., Las Casas, D. de, Hendricks, L. A., Welbl, J., Clark, A., Hennigan, T., Noland, E., Millican, K., Driessche, G. van den, Damoc, B., Guy, A., Osindero, S., Simonyan, K., Elsen, E., Sifre, L. (2022). *Training compute-optimal large language models*. https://arxiv.org/abs/2203.15556

Kaplan, J., McCandlish, S., Henighan, T., Brown, T. B., Chess, B., Child, R., Gray, S., Radford, A., Wu, J., & Amodei, D. (2020). *Scaling laws for neural language models*. https://arxiv.org/abs/2001.08361