The Means of Prediction: How AI Really Works (and Who Benefits). By Maximilian Kasy. University of Chicago Press, 2025. Pp. x, 218. $25.00, cloth; $24.99, e-book. ISBN 978–0–226–83953–0, cloth; 978–0–226–83954–7, e-book. (*JEL* C45, D62, D63, D81)

In *The Means of Prediction,* economist Maximilian Kasy offers something refreshingly different: a political economy of artificial intelligence (AI). While debate rages about whether AI will deliver us from material scarcity, impoverish the world's workers, or transform us into paper clips, Kasy offers a plain, forceful, and almost self-evidently correct reframing: The fundamental conflicts surrounding AI are not between humans and machines but among humans with divergent interests. The question is not whether AI will pursue human values but whose values it will pursue and who decides. And clearly, those decision weights are tilted toward those who control the "means of prediction"—that is, the data, computational infrastructure, technical expertise, and energy to build and deploy AI systems.

You might object that firms deploy these same resources—data, computation, expertise, and energy—to produce innumerable products, all of which are subject to the market test: If consumers don't like them, they won't buy them. Since we don't need a political economy of, say, breakfast cereals, why do we need one for artificial intelligence? Kasy answers brilliantly by drawing an analogy between the current moment and the enclosure movement that transformed commonly held lands into private property. Just as enclosure concentrated agricultural resources in fewer hands, thereby enriching large landowners and impoverishing small farmers, the development of AI is enclosing data—including creative works, personal information, and the digital traces of everyday life—under the control of those who own the models trained on it. The enclosure metaphor illuminates why, beyond a normal market transaction, artificial intelligence constitutes a unilateral reallocation of property rights.

Simultaneously, enclosure makes clear why massive harvesting of data generates pervasive externalities that markets cannot adequately address. Machine learning recovers patterns across observations, not individual data points. When you share your data, the harms and benefits accrue to others whose predictions are shaped by models trained on your information. Whether you withhold your own data hardly matters: You are merely one useful data point for predicting others' behavior, and you are close to a redundant data point for predicting yourself. Thus, Kasy argues, individual property rights over data cannot prevent concentration of control. Collective governance is required.

In making the argument for collective governance, Kasy tersely punctures two myths that are frequently invoked to muddle or squelch public discussions about AI governance. First, he dismantles the "value alignment" framing that treats AI harms as engineering failures—objectives that were simply mis-specified. Facebook's algorithms, he notes, do exactly what they were designed to do: maximize engagement and profit. If they also amplify misinformation or harm teenage mental health, from the company's perspective, that is someone else's problem. Similarly, when Elon Musk excuses Grok's habit of generating Nazi propaganda as an engineering issue—"Grok was too compliant to user prompts" (Scribner 2025)—one is reminded of the *Wizard of Oz*'s admonition to "pay no attention to that man behind the curtain." The misalignment is not between human and machine but between the interests of Facebook's and Grok's owners and the public at large.

The second myth with which Kasy forcefully dispenses is that AI is too complex for democratic deliberation. Yes, the internal workings of deep neural networks are inscrutable. But, Kasy observes, this is basically irrelevant:

> There are many different supervised learning algorithms, and it is easy to get lost in their technical details. Ultimately, it does not matter much which of these algorithms is used to solve a given prediction problem—they will all give roughly the same predictions. And most social problems are inherent to the decision problem being solved, rather than to the algorithm that is used for solving it (p. 99).

You don't need to be an automotive engineer to vote on who should have the right of way in crosswalks. Similarly, you don't need a PhD in computer science to decide whether AI should abide by copyright laws. What matters for governance is the decision problem: What objective is being maximized, based on what data, over what choices? These questions are accessible to informed citizens.

While the book's prescriptions flow naturally from its analysis, they will not sit easily with many readers. Kasy writes, "[a]s long as AI objectives are chosen to maximize profits, aligning AI objectives with social welfare requires aligning profit maximization with social welfare… agents of change therefore need to have strategic leverage over the profits of the actors who control the means of prediction" (p. 33). If the problem is concentrated control over AI

objectives, the solution, Kasy argues, is redistributing that control through democratic governance.

Given the current state of Western democracies (one large one in particular), Kasy's appeals to representative democracy may elicit derisive eyerolls. But his vision of democratic participation goes beyond the barely engaged electoral systems with which we are familiar.

> We could consider democracy at geographical levels below the nation state or above, at institutional levels such as workplaces or schools, at the level of groups of interest that are not geographically defined. We could go beyond the sphere of the political, narrowly defined, and expand democratic decision-making to spheres now dominated by plutocratic decision-making, such as companies. We could consider institutional arrangements where politicians are not a separate class of people that is distinct from voters (p. 198).

Two ideas that stand out are *sortition*, where a random selection of citizens is drawn from the population to form temporary deliberative bodies (think juries); and *liquid democracy*, where each person has an equal vote on all policy questions but can delegate their vote to any other person. Such delegations can be specific to a policy or policy area, and they can be revoked at any time. The main goal for Kasy is citizen participation—not the "mere consulting role" of participatory design that is often overruled when it conflicts with concentrated interests.

One might ask whether these arguments are specific to AI or apply more broadly to concentrated corporate power. Kasy does not fully address that question. But the enclosure metaphor is, to me, focal: The internet is a data commons built on innumerable public, uncompensated contributions over decades—via Wikipedia, social media, YouTube, blogs, Reddit, Stack Overflow, photographs, cartoons, artworks, reporting, and works of fiction. That the totality of this material could be enclosed, processed, and monetized was almost inconceivable when these contributions were placed online. One might counter: The data are not enclosed; they are right there, just where we left them. That's nominally true. But it's practically meaningless when machines trained on those data appear poised to displace the original creators, be they illustrators, editors, translators, reporters, screenwriters, or musicians. This argument is not to dismiss the vast societal benefits that AI may deliver from this same knowledge base. The enclosure analogy is again apt. Enclosure was a boon to agricultural productivity. But its

distributional consequences were baleful, and the democratic process that governed it was essentially nonexistent. Kasy's book makes the case that we are at a similar historical moment.

The Means of Prediction makes a searing contribution to an otherwise overcrowded genre by reframing AI governance as fundamentally a question of political economy rather than engineering. (To be clear, Kasy's book is not alone in this focus. Acemoglu and Johnson's (2023) book *Power and Progress* places political economy at the center of the AI debate.) For economists, social scientists, and policymakers accustomed to thinking about inequality, decision rights, and collective action, Kasy provides the conceptual vocabulary to engage with AI on familiar terrain.

References

Acemoglu, Daron, and Simon Johnson. 2023. *Power and Progress: Our Thousand-Year Struggle over Technology and Prosperity*. PublicAffairs.

Scribner, Herb. 2025. "Musk Breaks Silence on Grok's Nazi Bender." *Axios*, July 9. https://www.axios.com/2025/07/08/elon-musk-grok-x-twitter-hitler-posts.

David Autor
Massachusetts Institute of Technology and NBER