

Brookings Papers

ON ECONOMIC ACTIVITY

BPEA Conference Draft, September 24-25, 2026

Why Is AI So Contentious?

Martin Beraja (University of California, Berkeley)

Noam Yuchtman (London School of Economics)

The Brookings Institution is committed to quality, independence, and impact. We are supported by [a diverse array of funders](#). In line with our [values and policies](#), each Brookings publication represents the sole views of its author(s).

Why is AI so Contentious?

Martin Beraja (UC Berkeley) and Noam Yuchtman (LSE)*

Prepared for the Brookings Papers on Economic Activity

September 9th, 2026

1. A Contentious Technology

Artificial intelligence (AI) has been heralded by innovators, policymakers, and academics as a technology with enormous potential contributions to economic growth — some models predict that GDP could double (Beraja and Talamás, 2026) or even grow much beyond that (Aghion et al., 2019; Davidson, 2021; Jones, 2026).¹ It is also a technology in which large numbers of people have a financial stake: tens of millions of people have either direct or pension fund investments in the S&P 500, around a third of which (by value) is currently comprised of “magnificent seven” companies (Alphabet, Amazon, Apple, Meta, Microsoft, Nvidia, and Tesla), all closely linked to AI development.

Despite AI’s transformative potential, and direct economic interest, the technology is far from uniformly lauded. Large numbers of people have a negative outlook on AI. Survey results from the Pew Research Center (Figure 1, Panel A) indicate that around half of American adults are more concerned than excited about increased use of AI in daily life.² Only 10% are more excited than concerned, and this number has fallen by nearly half over the last five years. Survey results from the Ada Lovelace and Alan Turing Institutes (Panels B and C) show that although 63% of UK adults view large language models as beneficial, 47% of them view them as concerning.³

What are people concerned about? OpenAI’s own assessment is a good place to start, both for what it contains and for what it omits:

*Martin Beraja: maberaja@berkeley.edu. Noam Yuchtman: n.yuchtman@lse.ac.uk.

¹See also Nordhaus (2021) for a more cautious analysis.

²The data and original report can be found at: <https://www.pewresearch.org/science/2025/09/17/how-americans-view-ai-and-its-impact-on-people-and-society/>, last accessed August 3, 2026.

³The data and original report can be found at: <https://attitudestoai.uk/findings-2025/benefits-and-concerns>, last accessed August 3, 2026.

*While we strongly believe that AI’s benefits will far outweigh its challenges, we are clear-eyed about the risks—of jobs and entire industries being disrupted; bad actors misusing the technology; misaligned systems evading human control; governments or institutions deploying AI in ways that undermine democratic values; and power and wealth becoming more concentrated instead of more widely shared.*⁴

The quote speaks of the risk that AI will disrupt jobs and concentrate wealth, and the risk that it will be misused or escape control. These are two kinds of concerns, and sources of conflict, that a transformative technology can generate: *economic* and *existential*. Both, we will argue, are familiar: we have seen major technologies raise these concerns in the past. The quote omits, however, how the technology was built: by taking vast quantities of human-generated content. We will argue that this appropriation had an economic logic, the same one that has driven innovation at the resource frontier throughout history (Beraja and Yuchtman, 2026), and that allowing some of it may even be desirable. But it has raised a third kind of concern, a *moral* one, with no clear counterpart in earlier technologies. This moral concern, not the return of familiar ones, is why AI is *so* contentious and the conflict *so* visceral.

2. Economic and Existential Concerns

The first source of conflict is *economic*. AI threatens to displace workers across a wide range of occupations — from drivers and warehouse workers to paralegals, programmers, and other workers in the knowledge economy — and to do so faster than workers and institutions can adjust (Brynjolfsson et al., 2025; Restrepo, 2025). It also threatens to concentrate the gains from the technology in very few hands: Forbes identified 45 “AI Newcomers” to its list of billionaires in 2026.⁵


These concerns are real, but they are not new. Every major technological transition has displaced workers and enriched a few innovators, and economists have long studied both the resulting disruption and the policies that might address it. In our own work, we have shown that technological transitions can happen fast or slow depending on how specific workers’ skills are (Adão et al., 2024); and that slowing down automation can be desirable when firms do not internalize the losses of displaced workers who are financially vulnerable (Beraja and Zorzi, 2024). Of course, the intensity of the economic disruption from AI may be different

⁴From OpenAI’s 2026 report, “Industrial Policy for the Intelligence Age: Ideas to Keep People First,” available at: <https://cdn.openai.com/pdf/561e7512-253e-424b-9734-ef4098440601/Industrial%20Policy%20for%20the%20Intelligence%20Age.pdf>, last accessed August 3, 2026.

⁵See: <https://www.forbes.com/sites/phoebeliu/2026/03/10/meet-the-45-ai-newcomers-to-forbes-2026-billionaires-list/>, last accessed August 4, 2026.

Figure 1: Negative assessments of AI in the US and UK

Panel A: Attitudes towards AI in the US




Source: Pew Research Center

Panel B: Perceptions of AI Benefits in the UK

Figure 3: The extent to which each AI use is perceived as beneficial

(Due to rounding, percentages may not total 100%)

⑦ "To what extent do you think that the use of this technology will be beneficial?"




Source: Ada Lovelace and Alan Turing Institutes

Panel C: Perceptions of AI Concerns in the UK

Figure 4: The extent to which each AI use is perceived as concerning

(Due to rounding, percentages may not total 100%)

⑦ "To what extent are you concerned about this use of this technology?"



Source: Ada Lovelace and Alan Turing Institutes

from technologies of the past: in [Beraja and Yuchtman \(2025\)](#) we have discussed how AI may disrupt the very institutions — labor markets, tax systems, and property rights — through which societies have historically cushioned technological change. Nor is concern about concentrated wealth new: the railroad and oil fortunes of the Gilded Age produced their own “robber barons,” and their own political backlash.

Indeed, the economic concerns raised by AI go back at least to the industrial revolution. The Church’s first social encyclical, Leo XIII’s *Rerum Novarum* ([Leo XIII, 1891](#)), was written in response to “the vast expansion of industrial pursuits and the marvellous discoveries of science,” and its subject was distribution: “the enormous fortunes of some few individuals, and the utter poverty of the masses,” and workers “surrendered, isolated and helpless, to the hardheartedness of employers and the greed of unchecked competition.”⁶

The second source of conflict is *existential*. AI may be misused by bad actors to design pathogens or mount cyberattacks; it may be embedded in autonomous weapons; and, some fear, sufficiently capable systems may escape human control altogether. These are concerns not about who gains from the technology but about whether humanity survives it. They, too, are not new. Nuclear technology raised precisely this kind of concern, and it generated its own encyclical: John XXIII’s *Pacem in Terris* ([John XXIII, 1963](#)) described people “living in the grip of constant fear” that a nuclear “conflagration could be started by some chance and unforeseen circumstance,” and warned that “the very testing of nuclear devices for war purposes can, if continued, lead to serious danger for various forms of life on earth.”⁷ Its remedy was blunt: “Nuclear weapons must be banned.”

3. A New Moral Concern over Appropriation

3.1. *How the technology was built*

The developers of large language models exploited an extraordinary insight (complementing their technical innovations): the masses of online text that had been posted and archived over decades since the emergence of the world wide web could be used as *training data* for artificial intelligence models. Acting on this insight, they assembled the training corpora of the leading models by scraping essentially the entire accessible web: Common Crawl, Wikipedia, Reddit, news sites, personal blogs, forums, and, as litigation later revealed, online libraries of pirated books.

⁶*Rerum Novarum*, paragraphs 1 and 3. Full text available at: https://www.vatican.va/content/leo-xiii/en/encyclicals/documents/hf_l-xiii_enc_15051891_rerum-novarum.html, last accessed September 2, 2026.

⁷*Pacem in Terris*, paragraph 111. Full text available at: https://www.vatican.va/content/john-xxiii/en/encyclicals/documents/hf_j-xxiii_enc_11041963_pacem.html, last accessed September 2, 2026.

What was taken was, first, people’s own *words*: the questions they asked and answered on forums, the entries they wrote and edited, the opinions they posted, the journalism they reported. It was, second, people’s own *creations*: their books, their photographs, their paintings and illustrations, their music, their code. And it is now, increasingly, people’s own *expertise*: what they know and do at work. Uber’s drivers are being used to train the autonomous driving systems meant to replace them;⁸ Meta has asked its own employees to record their work to train models, before facing resistance;⁹ firms such as Mercor are paying experts to teach models to do their jobs.

We see this process of innovation intertwined with appropriation as the latest instance in a long history of innovations that expanded the frontier of resources available for economic use, and whose economic logic we have studied in [Beraja and Yuchtman \(2026\)](#).¹⁰ Mining and metallurgy brought previously unusable minerals into use; drilling gave access to underground fossil fuels; radio turned the electromagnetic spectrum into a medium for transmitting information. A distinctive feature of resources outside the frontier is that property rights — who owns them and how they can be used — are typically poorly defined. Before the innovations that revealed their value, it had not been worthwhile, or even possible, to establish ownership and specify rights over their potential uses (as in [Demsetz, 1974](#)). This gives innovators an opportunity to claim the newly valuable resource, and in important historical cases they did so, in anticipation of extraordinary rents. But there are often counterparties with a claim on the resource, sometimes the state by default, and once they become aware of its value they contest the innovator’s claim. Appropriation, in this sense, is not an aberration of innovation at the frontier but part of how it proceeds, and, as we argue in the conclusion, some of it may be desirable.

Human-generated data fit this pattern exactly. Before the insight that they could serve as training data, the words people had posted online, the works they had created, and the ways they did their jobs were abundant, cheap (or free), and largely ignored as resources. For online content, property rights were oriented around copyright restrictions that allowed the “fair use” of content when it was transformative and used for purposes of education and research. AI developers felt that their use of online content to train their models satisfied

⁸See <https://techcrunch.com/2026/05/01/uber-wants-to-turn-its-millions-of-drivers-into-a-sensor-grid-for-self-driving-companies/>, last accessed August 4, 2026.

⁹See: <https://medium.com/loop-of-thought/meta-just-made-its-employees-train-their-own-replacement-46d52365bd72>, last accessed August 4, 2026.

¹⁰We do not argue that the resource-frontier lens accounts for the economic or existential concerns discussed above; those concerns are real and are shared, in different measure, by past technologies. Our claim is that the lens accounts for what is distinctive about the conflict over AI. For other important critical perspectives on AI innovation, see [Acemoglu \(2022\)](#) and [Carvalho \(2025\)](#); in some of our own work we have also highlighted concerns regarding the alignment of AI with the aims of autocrats ([Beraja et al., 2023, 2025](#)).

this criterion.¹¹ There was no market for training rights, no established price to acquire data for training, and no shared understanding of what it would even mean to “license” language or images for this purpose. Over other forms of human-generated data, such as how a driver drives or how a lawyer drafts a brief, no rights had been specified at all.

One could imagine a counterfactual world in which, rather than appropriate human-generated data, companies attempted to contract for them, with publishers, platforms, and workers.¹² This might have avoided subsequent conflict, but it would not have been straightforward to implement: innovators would have been offering to pay to use resources that did not yet have a market; they would either be forced to divulge details of their innovation (“tipping their hand” to the counterparty) or they would be offering to pay for data in a way that would appear extremely opaque to the other side. Moreover, such a (challenging) contract would need to be negotiated repeatedly with many content producers and, in the case of expertise, with millions of individual workers. Facing such contracting frictions, innovators decided to appropriate without contracts. None of it was done with consent. The people whose words, creations, and expertise became training data were not asked, were not paid, and in most cases did not know.

Conflict followed, exactly as it had in past instance on innovations at the resource frontier. Once the value of human-generated data *for AI training* became clear with the launch of ChatGPT, counterparties with a claim on those data contested the free access that companies like OpenAI enjoyed. In December 2023, *The New York Times* sued OpenAI and Microsoft, alleging that its articles had been used without permission violating copyright. Similar lawsuits regarding the fair use of text and images by generative AI companies have since been filed by a diverse set of counterparties, including actors, movie studios, traditional reference sources (e.g., Encyclopedia Britannica and Merriam-Webster), novelists, and more. As of this writing, multiple lawsuits are ongoing, with a large, consolidated case representing multiple plaintiffs in multiple US judicial districts recently transferred to the Southern District of New York.¹³ There has been one notable settlement: the *Bartz v. Anthropic* case produced the largest copyright settlement in history, with Anthropic agreeing to pay \$1.5 billion to a group of authors and publishers.¹⁴ Conflict over the appropriation of expertise emerged

¹¹See, for example, OpenAI’s statement stating explicitly, “Training is fair use, but we provide an opt-out because it’s the right thing to do”. Online at: <https://openai.com/index/openai-and-journalism/>, last accessed August 4, 2026.

¹²There was some effort at contracting to purchase text. For example, Facebook considered an acquisition of publisher Simon & Schuster precisely to acquire training data. See: <https://www.theguardian.com/books/2024/apr/09/meta-discussed-buying-publisher-simon-schuster-to-train-ai>, last accessed August 4, 2026.

¹³See: https://www.bloomberglaw.com/public/desktop/document/INREOpenAIIncCopyrightInfringementLitigationDocketNo3143USJPMLDec/3?doc_id=X10L562HT9O9VHRKA8D1VE5I4TK, last accessed August 4, 2026.

¹⁴It is worth noting that the primary legal basis for Anthropic’s settlement in this case was not violations of

later and has so far taken the form of worker resistance rather than litigation, as at Meta, but there is no reason to expect it to stay that way.

The contest is not only judicial. Appropriation must ultimately be ratified by the state, and so a second prediction of the resource frontier framework is that appropriators invest in cultivating political influence. The state’s willingness to ratify the appropriation of human-generated data is inevitably a matter of politics as well as law, and, to some extent, a function of companies’ lobbying efforts and political connections. This echoes an earlier wave of appropriation of user data, which were converted into predictive inputs used for online marketing by “surveillance capitalists” like Google and Facebook (Zuboff, 2019).¹⁵ Indeed, AI companies are among the most active lobbyists in both the US and in Europe, and their activity is growing. Between 2020 and 2026, quarterly US federal lobbying spending by major technology firms has nearly doubled, from around \$10 million to \$20 million (see Figure 2, Panel A). Among the technology firms, the largest growth (in percentage terms) in lobbying spending in 2026 was exhibited by the LLM developers Anthropic, where lobbying spending *quadrupled* between 2025 and 2026, and OpenAI, where lobbying spending nearly doubled (see Figure 2, Panel B).

In the EU, lobbying by technology companies is also intense — the Corporate Europe Observatory ranks the technology sector as the largest spender, above energy and finance.¹⁶ In Europe, the key counterparty engaged in conflict with technology companies is “the state” as a whole (representing populations of individuals concerned about privacy, ethics, etc.). Technology company lobbying aims to relax the EU’s data regulations, which arguably stifle AI innovation on the continent. Efforts by AI firms to influence policy naturally embed them in the contentious political arena, adding an important dimension to the broader social debates over the technology.

3.2. *The Original Sin*

The appropriation of human-generated data is, we argue, the root of a third kind of concern, a *moral* one, and the reason AI is so contentious.

In Beraja and Yuchtman (2026) we document many cases in which innovators appropriate fair use generically, but specifically their use of training data on websites possessing pirated copies of books. See: <https://legalblogs.wolterskluwer.com/copyright-blog/the-bartz-v-anthropic-settlement-understanding-americas-largest-copyright-settlement/>, last accessed August 4, 2026.


¹⁵The conflict over appropriated user data has also been extremely high-stakes. For example, Facebook settled a lawsuit over data sharing with a third-party (Cambridge Analytica) for \$725 million. See: <https://www.theguardian.com/business/2022/dec/23/facebook-owner-to-settle-class-action-suit-over-cambridge-analytica-scandal>, last accessed August 4, 2026.

¹⁶See <https://corporateeurope.org/en/2026/06/eu-corporate-lobby-league-2026>, last accessed August 4, 2026.

Figure 2: Lobbying by technology firms

Panel A: Quarterly spending on US federal lobbying, 2020–2026

Aggregate first-quarter spending by major tech, social media, and AI players from 2020 through 2026




Note: Meta is the parent company of Instagram and Facebook, Alphabet is the parent company of Google and Youtube, Anthropic is the parent company of Claude.ai, and OpenAI is the parent company of ChatGPT
Source: Issue One review of federal lobbying reports

Source: Issue One

Panel B: Growth in US federal lobbying spending, 2025–2026

Comparative first-quarter spending by Meta, Alphabet, Microsoft, Nvidia, Anthropic, and OpenAI between 2025 and 2026



Note: Meta is the parent company of Instagram and Facebook, Alphabet is the parent company of Google and Youtube, Anthropic is the parent company of Claude.ai, and OpenAI is the parent company of ChatGPT
Source: Issue One review of federal lobbying reports

Source: Issue One

priated resources outside the frontier — minerals, oil, the electromagnetic spectrum — and where several of the resulting conflicts were intense. Yet those conflicts were fought over revenue streams and legal property rights. Thus, appropriation partly explains why there is conflict over AI; it does not explain why the concerns underpinning conflict are moral and the conflict itself visceral.

That the concerns are moral and the conflict visceral is evident in the vocabulary used by the aggrieved. The authors, journalists, photographers, artists, and media corporations that have brought high-stakes legal cases against AI developers¹⁷ do not claim only that they were not paid. Their claim is that AI was *built from* them, without permission. The vocabulary of this conflict is one of theft and violation rather than of adjustment and compensation. Past technologies certainly attracted moral objections, but these concerned their *uses* — whether a weapon should be built or an embryo edited. The moral objection to AI is about its *inputs*. Tellingly, the Church’s response to AI focuses precisely on the moral violation of human sanctity. Leo XIV’s *Magnifica Humanitas* (Leo XIV, 2026), “On Safeguarding the Human Person in the Time of Artificial Intelligence,” is less about the distribution of gains, or about survival, than about the danger of translating “the mystery of the person, into data and performance,” and about the duty to “remain profoundly human” in the face of a technology whose imitation of us threatens “new forms of dehumanization.”¹⁸

The reason the conflict is so visceral thus lies in the *nature* of the resource appropriated, not only in the fact of appropriation. Land, oil, and the spectrum can be bought, sold, seized, and, once property rights are settled, paid for. There was never any doubt that they could be taken. Human creativity, imagination, and expertise were thought to be different. They are what people make and know, and to many an expression of who they are — to some, their very essence. Human-produced works could be copied, of course: books have always been pirated, and pictures and films have always leaked. But a pirated book is still that author’s book, and the author’s capacity to write is not taken by the pirate. Copying deprived creators of payment, not of true authorship; the work remained theirs, and the ability to produce it remained theirs alone too. What large language models did is not copying. They absorbed millions of works into a single system that retains no trace of who made what, and that learned from them to *do what the makers do* — and, in a sense, to be something like what they are. What was taken, then, was not a copy of the work but

¹⁷See <https://www.debevoise.com/insights/publications/2025/12/ai-intellectual-property-disputes-the-year-in>, last accessed August 4, 2026.

¹⁸*Magnifica Humanitas*, paragraphs 10 and 15. Full text available at: <http://www.vatican.va/content/leo-xiv/en/encyclicals/documents/20260515-magnifica-humanitas.html>, last accessed September 2, 2026. Notably, the encyclical is dated on the 135th anniversary of *Rerum Novarum*, and Leo XIV has explained his choice of name by reference to Leo XIII’s response to the first industrial revolution.

something else: the authorship of what one has made, and with it the belief that to write, to draw, and to know were things only a human could do. That is what people mean when they say their work is being stolen, even where no law was broken.

The concern therefore goes beyond economics, to the diminution of what it is to be human and to create. The cartoonist Sarah Andersen wrote about this more eloquently than we can, in the *New York Times*:

I was sent via Twitter an image generated by A.I. from a random fan who had used my name as a prompt. It wasn't perfect, but the contours of my style were there. ... I felt violated. The way I draw is the complex culmination of my education, the comics I devoured as a child and the many small choices that make up the sum of my life. The details are often more personal than people realize—the striped shirt my character wears, for instance, is a direct nod to the protagonist of “Calvin and Hobbes,” my favorite newspaper comic. ... Art is deeply personal, and A.I. had just erased the humanity from it by reducing my life’s work to an algorithm.

A taking of this kind, following a breakthrough of this magnitude, is AI’s “original sin.” We choose the term with care. In Genesis, the first transgression was the taking of the fruit of the tree of knowledge. It was a wrong; yet the knowledge it brought made everything that followed possible, and a long tradition has called the fall, for that reason, fortunate. So it may be with AI. The insight was brilliant, the breakthrough remarkable, and it promises to bear a great deal more fruit. But it was humans, not God, who were aggrieved.

4. Conclusion

We have argued that AI is so contentious — and the conflict over it so visceral — because, on top of the economic and existential concerns that every transformative technology raises, it raises a moral one: that it was built from people’s words, creations, and expertise, taken without consent. We have not argued that AI developers did something illegal. Whether they did is not clear to us, and it is not a question this essay can settle. The courts will decide whether training on copyrighted content is fair use — and reasonable people disagree about how they should rule. Similarly, bargaining and labor institutions will reach agreements about training on worker’s data.

What we can say is that the original sin was not incidental to the innovation. In [Beraja and Yuchtman \(2026\)](#), we show that poorly defined property rights at the resource frontier are not an obstacle to innovation but, in an important sense, a reason for it: the prospect

of claiming a resource before its value is recognized, and before anyone else has a claim on it, is part of what makes pushing the frontier worthwhile. Had online content been licensed before its training value was understood, there would have been nothing to license; and once its value was understood, the contracting frictions we described above would have made licensing at scale slow, costly, and perhaps impossible. Some degree of appropriation may therefore be the price of innovation at the frontier, and a society that forbids it altogether may find that the frontier stops moving. This is why, historically, the taking that followed the making has more often than not been legitimized after the fact.

This leaves societies confronting AI with two mistakes on offer. The first is to ignore the taking: to view the innovation as pristine, treat the grievance of human creators as confusion or rent-seeking, and be unprepared for the conflict that appropriation has always provoked. Many innovators, and their supporters, are making this mistake. The second is the opposite: to reject all appropriation, treat every contested resource as untouchable until the rules are settled, and so shut down the very innovations that push humanity to the next frontier. Many of AI’s critics are making this one. Each is half right. The innovators are right that the appropriation was inseparable from the breakthrough, and that a world which had waited for permission would likely still be waiting. The critics are right that what was taken belonged, in a sense no prior resource outside the frontier did, to the people it was taken from, and that no ruling on fair use will make the grievance disappear.

The resolution, as at every earlier frontier, will be worked out ex post, through lawsuits, licensing deals, legislation, and the new institutions that such conflicts have often produced. Our aim here has been narrower: to explain why the conflict is taking the form it does, and why this technology, unlike its predecessors, produced not only winners and losers but a moral grievance from its original sin. Understanding that is the first step to settling the conflict well. It means that the policies most often discussed, such as universal basic income, retraining programs, or pauses on AI development, will at best be incomplete because they address the economic and existential concerns, not the moral one. Addressing the grievance itself requires private or public solutions that restore to human creators some measure of ownership and attribution of what has been, and is still being, built from them.

References

- Acemoglu, D. (2022). Harms of AI. In Bullock, J. B. et al., editors, *The Oxford Handbook of AI Governance*. Oxford University Press.
- Adão, R., Beraja, M., and Pandalai-Nayar, N. (2024). Fast and slow technological transitions. *Journal of Political Economy Macroeconomics*.

- Aghion, P., Jones, B. F., and Jones, C. I. (2019). Artificial intelligence and economic growth. In Agrawal, A., Gans, J., and Goldfarb, A., editors, *The Economics of Artificial Intelligence: An Agenda*, pages 237–282. University of Chicago Press.
- Beraja, M., Kao, A., Yang, D. Y., and Yuchtman, N. (2023). AI-tocracy. *The Quarterly Journal of Economics*, 138(3):1349–1402.
- Beraja, M., Kao, A., Yang, D. Y., and Yuchtman, N. (2025). Exporting the surveillance state via trade in AI. Working paper.
- Beraja, M. and Talamás, E. (2026). The value of codifying organizational knowledge. Working paper.
- Beraja, M. and Yuchtman, N. (2025). Generalized disruption: Society, work, and property rights in the age of AI. In *The Political Economy of Artificial Intelligence*.
- Beraja, M. and Yuchtman, N. (2026). Growth at the resource frontier. Working paper.
- Beraja, M. and Zorzi, N. (2024). Inefficient automation. *The Review of Economic Studies*.
- Brynjolfsson, E., Korinek, A., and Agrawal, A. (2025). A research agenda for the economics of transformative AI. Working paper.
- Carvalho, J.-P. (2025). The political-economic risks of AI. Working paper.
- Davidson, T. (2021). Could advanced AI drive explosive economic growth? Open Philanthropy.
- Demsetz, H. (1974). Toward a theory of property rights. In *Classic Papers in Natural Resource Economics*, pages 163–177. Springer.
- John XXIII (1963). *Pacem in terris*. Encyclical letter, Vatican.
- Jones, C. I. (2026). A.i. and our economic future. Working paper.
- Leo XIII (1891). *Rerum novarum*. Encyclical letter, Vatican.
- Leo XIV (2026). *Magnifica humanitas*: On safeguarding the human person in the time of artificial intelligence. Encyclical letter, Vatican.
- Nordhaus, W. D. (2021). Are we approaching an economic singularity? Information technology and the future of economic growth. *American Economic Journal: Macroeconomics*, 13(1):299–332.
- Restrepo, P. (2025). We won’t be missed: Work and growth in the era of AGI. In *The Economics of Transformative AI*. NBER Chapters.
- Zuboff, S. (2019). *The Age of Surveillance Capitalism: The Fight for a Human Future at the New Frontier of Power*. PublicAffairs.