THE

DIGITALIST PAPERS

ESSAY

INFORMATION IN THE AGE OF AI:
CHALLENGES AND SOLUTIONS



INFORMATION IN THE AGE OF AI: CHALLENGES AND SOLUTIONS

JOSEPH E. STIGLITZ AND
MÀXIM VENTURA-BOLET

Transformative AI will transform how society produces and shares information. As Joseph Stiglitz and Màxim Ventura-Bolet argue, it could drive innovation by processing knowledge faster and more efficiently, even generating new questions and insights. Yet it may also erode the supply of reliable information, amplify mis- and disinformation, and leave corrective efforts underprovided. The authors present a framework to understand these informational challenges and outline practical policy responses.



I. Introduction

Artificial intelligence has been heralded as opening new frontiers in the accumulation of knowledge, as problems that would have taken years to solve—or were essentially insoluble—have been answered in short order. AI, especially transformative AI (TAI), has the potential to significantly increase the pace of innovation, not only processing information more efficiently and rapidly, but even asking and answering new questions.

The impact of AI on the information ecosystem lies at the heart of this transformation. How is AI reshaping the ways information flows through society? In what respect is it an improvement—and where do the dangers lie? Overall, are we becoming more informed as a society? Looking ahead, how might these effects evolve as technology advances? These are questions of immense consequence, but also of exceptional complexity.

Even before the advent of AI, there was a concern that digital platforms (search engines and social media) might lead to a deterioration in the quality of the overall information ecosystem. The downside risks are evident in the cesspool of mis- and disinformation that flourishes on the internet, with increasing concerns about the social, political, and economic consequences. There is the risk that AI might not only fail to correct the problems but also exacerbate them.

In this sense, we identify three separate problems, related to the digital platforms, but even more so to AI. The first is an undersupply of good information. Even with strong intellectual property laws, there are knowledge spillovers (externalities), and this is true not only for what we think of as "news," but also for other forms of information and knowledge. Those who read an informative newspaper article, based on expensive investigative research, relay that information to those they talk to, and these individuals thus obtain such information (if only part of it) for free. In a sense, all producers of information contribute, in

some measure or another, to the pool of knowledge that is available. Others take out from this pool, and neither those who use the pool of knowledge nor those who contribute to it pay or receive compensation commensurate with what they add or take out, and accordingly there is no presumption of optimality in the production, transmission, and usage of information.

The technology platforms and AI have, as we will explain below, made matters worse, because they are not just competitors of the traditional producers and disseminators of information; they are, in part, competitors, who take (or some say, steal) information from these traditional producers of information, use it to enhance their profits, simultaneously destroying the business model of traditional producers of information.

The second problem is that there is an oversupply of mis- and disinformation. Those who thereby pollute the information ecosystem don't pay for the consequences, the damage that occurs as a result of a poorer information ecosystem, and the resources that have to be deployed to separate out truthful information from the mis- and disinformation to which they have contributed. Indeed, in some cases, they profit financially from such pollution. Some technology platforms have a business model of "engagement through enragement," polarizing society as they let algorithms amplify mis- and disinformation that riles up their users. It is hard not to believe that some AI companies won't act similarly, spewing forth ever better summaries of news and other information tailored to the preexisting biases of their users. Moreover, some firms have a political agenda, and use their technology and algorithms to promote it, without regard to the truthfulness of the information that they disseminate. Finally, mis- and disinformation is cheaper to provide than high-quality information, and there is an ample

Mis- and disinformation is cheaper to provide than high-quality information, and there is an ample supply of "information providers" willing and able to exploit those unable to distinguish between the two.

supply of "information providers" willing and able to exploit those unable to distinguish between the two.

Just as there is an undersupply of truthful information and an oversupply of distortionary information, there is an undersupply of efforts to "correct" mis- and disinformation, for that too is a public good, from which others benefit. This is the third problem.

One view is that, despite these concerns, the institutional machinery of regulation cannot artfully guide the development of a technology as uncertain, fast-paced, and consequential as AI—and that attempts to do so will slow the pace of innovation, and accordingly societal welfare. In this perspective, the best course of action is to give free rein to market competition and hope that society will adapt naturally.[1]

The central thesis of this paper is that that view is totally misguided. It simply assumes that all innovation is welfare enhancing, when it should be obvious that that is not the case: Innovations in creating and exploiting market power are welfare decreasing. We show that there is a real risk that these technological innovations too may be welfare decreasing—in the absence of appropriate regulation. Here, we focus on only one set of harms: The threats to the information ecosystem are real. We show that these advances that have the power to improve the information ecosystem may have just the opposite effect.

Here, we focus on only one set of harms: The threats to the information ecosystem are real. We show that these advances that have the power to improve the information ecosystem may have just the opposite effect.

However, with the right legal and regulatory framework, there is a much greater chance that innovation in these areas will be beneficial to society as a whole, as opposed to just increasing the profits of the tech companies.

II. The Framework

We develop an analytical framework that distinguishes the channels through which AI enhances the production and dissemination of high-quality information from those through which it undermines them, and we use this framework to identify public-policy interventions capable of mitigating the adverse effects. This section builds on our 2025 paper, "The Impact of AI and Digital Platforms on the Information Ecosystem."[2] Here we distill the key ingredients of the model for a broader policy audience. Readers interested in the formal derivations and technical details can refer to the original paper.

The pre-AI information ecosystem

Understanding how AI reshapes the information ecosystem requires first recalling how that system functioned—imperfectly but coherently—before the arrival of digital platforms and AI. Then, producers of new information ("news")—newspapers, academic journals, broadcasters, and later online media—created content that contributed to society's stock of knowledge. Information has strong public-good characteristics: Once produced, it can be used ("consumed") by anyone without depletion (technically, economists say consumption is "non-rivalrous"). The critical property of a public good is that its social value exceeds what producers can capture privately, which is why if there is to be an adequate supply, there must be public support.

Information has strong public-good characteristics: Once produced, it can be used ("consumed") by anyone without depletion (technically, economists say consumption is "non-rivalrous").

Business models relied on attracting visits and attention, which could be monetized through advertising, subscriptions, or data collection. Consumers, for their part, accessed this information often at little or no direct cost.

This architecture balanced, however imperfectly, the public-good nature of information[3] (Arrow (1962), Stiglitz (1975, 1986, 1999, 2021)) with incentives for its creation and dissemination. And for much of the 20th century, producers largely operated under professional and reputational norms that favored truthfulness. There was never a presumption that it yielded an optimal quantity or composition of information, but it provided a workable equilibrium. In some places, this equilibrium was already fraying before the rise of digital platforms and AI. For instance, over time, some also found economic and strategic incentives to generate untruthful, biased, or sensational content. Now, our news information ecosystem is at risk of collapse.

Motivation and the effects of AI

Against this backdrop, our inquiry begins with a simple question: How will AI affect the information ecosystem? AI introduces several powerful and interrelated shocks:

1. **Improved production and dissemination**. AI systems dramatically enhance the efficiency of processing and transmitting information, allowing content to reach broader audiences faster and at lower cost.
2. **Erosion of the producer business model**. By providing answers and summaries directly to users, AI intermediaries reduce visits to original producers, weakening their revenue base.[4]
3. **Falling relative cost of lies**. Generative models make the creation of persuasive misinformation inexpensive compared with producing verified truth.
4. **The "drone-war" effect**. AI simultaneously enhances the ability to detect falsehoods and the ability to evade detection, leading to an arms race between verification and deception.

We develop a tractable model of the information ecosystem

that captures these forces. The model helps us understand how AI influences the equilibrium quantity of truthful information, the share of misinformation, and the degree of polarization in society.

Structure of the model and equilibrium outcomes

In our paper, we formalize these dynamics in a setting with two types of agents—producers and consumers—and analyze how their interactions determine what kind of equilibrium society reaches. There can be a truthful equilibrium, where mis- and disinformation is not produced; a mixed truth-lies equilibrium, where both co-exist; or an equilibrium where the information ecosystem collapses, where incentives to produce at least certain kinds of socially desirable information are so weak that such information is not produced at all.

The model makes several simplifications, both for analytic simplicity and to bring out starkly the risks presented by AI. It assumes, for instance, that news producers create either "truthful" or "untruthful" information and earn revenue from the attention they attract. Producing truth is more expensive than producing untruthful information. Consumers allocate their attention proportionally across producers. A share of consumers can recognize truth ("informed"), while the rest cannot ("uninformed").

Producing truth is more expensive than producing untruthful information. Consumers allocate their attention proportionally across producers.

Digital platforms and AI systems operate as intermediaries that shape both sides of the market: They determine how information is transmitted, how visible producers are, and how much of the economic value of attention accrues to intermediaries rather than to the creators themselves.

The framework highlights four key parameters: the efficiency of transmission (how easily information circulates), the degree of intermediation (the extent to which the information is processed, e.g.,

combining information from different sources), the ability of consumers to screen for truth, and the relative cost of lies. Their interaction determines which equilibrium the system converges to.[5]

A truthful equilibrium prevails when several reinforcing conditions hold:

- There is strong demand from informed consumers, meaning a high share of users who value and can identify accurate information.
- Demand from uninformed consumers is relatively weak, reducing incentives to serve low-information audiences with cheap misinformation.
- The cost of producing lies is high—either because deception is easily detected or because legal and reputational penalties make it expensive.

Under these conditions, producers have little reason to generate misinformation, and the flow of truthful content is robust. In such an environment, the overall quantity of information produced is highest when:

- Consumers engage directly with producers, so a larger share of attention translates into revenue for those who create new information (low capture of attention by intermediaries).
- The perceived value of information is high—either because it is socially important or because informed and uninformed consumers alike seek accuracy.
- The cost of producing information remains moderate, so that investing in verification and original reporting is profitable.

By contrast, when these conditions weaken, the economy moves toward a regime in which both truthful and false information coexist. This outcome is characterized by stabilizing dynamics: If the share of misinformation becomes too high, truthful information grows scarce and thus valuable, encouraging producers to shift back toward "truth."

If it is too low, misinformation becomes profitable again, drawing more producers into deception.

These equilibria capture the competing forces shaping today's information landscape. AI and digital platforms can amplify both dynamics—strengthening truthful dissemination under the right incentives or entrenching misinformation when market and institutional conditions favor it.

If the share of misinformation becomes too high, truthful information grows scarce and thus valuable, encouraging producers to shift back toward "truth." If it is too low, misinformation becomes profitable again, drawing more producers into deception.

Risks and limiting cases

The model also clarifies the risks that emerge when any of the key parameters move toward their extremes. Each has intuitive implications.

1. When intermediation (obtaining news through technology platforms or AI) becomes too high, most information consumption does not occur directly from producers. Even if these intermediaries efficiently summarize or distribute content, the original creators capture little of the economic value. As a result, incentives to invest in costly verification and original reporting collapse. The ecosystem may appear highly efficient from a user's perspective—information everywhere, instantly available—but beneath the surface, the production of new and accurate information dries up.

   This dynamic is analogous to the Grossman–Stiglitz paradox in financial economics: In a perfectly informationally efficient market, no one has an incentive to gather costly new information, because prices already reflect everything that is known.[6] Thus, a market in which prices perfectly transmitted information would be a market without information—only information that was absolutely costless to obtain would be reflected in the prices. Such a market would not function well. Similarly, if AI and platforms make existing information perfectly accessible to consumers, incentives to generate new truths vanish.

At the current state of AI, hallucinations and inaccuracies still serve as a natural brake on user substitution away from primary sources. Because users cannot fully rely on AI outputs, they continue to consult original information producers, thereby preserving—at least partially—the incentives for content creation. If AI systems become perfectly reliable—or reliable enough—and all information consumption is mediated through them, incentives to produce new information could collapse entirely, and so too would the quality of the information ecosystem.

2. When the cost of lies becomes too low, the economics of content production tilt toward misinformation. Cheap generative tools make it easy to flood the ecosystem with plausible but false narratives, while truthful reporting remains expensive and slow. The marginal producer—whether a news outlet or an influencer—finds it more profitable to pursue engagement rather than accuracy. The equilibrium share of misinformation rises, trust erodes, and the collective quality of knowledge deteriorates.
3. Low screening capacity means consumers cannot reliably distinguish true from false information. Demand for quality weakens, and misinformation becomes self-reinforcing: As trust decreases (they know that the quality of information has deteriorated), individuals are willing to pay less for news, further reducing producers' incentive to invest in it. The result is an "attention trap," where low-information consumers sustain high engagement with low-quality content, crowding out producers of truth.

Each of these limit cases illustrates a different path to information collapse. In the first, value capture shifts entirely to intermediaries; in the second, misinformation becomes too cheap; in the third, consumers lose the capacity or will to discriminate. None of these outcomes require TAI. They can—and in many contexts already do—arise from current systems that focus only on technologies related to distribution. But TAI could make each of these problems worse (particularly if TAI increases the difficulty of detecting mis- and disinformation faster than the ability to detect it.)

Why a Framework Matters

By reducing a sprawling debate to a clear structure, this framework clarifies what policymakers should target. Each of the four parameters—efficiency, intermediation, screening, and the cost of lies—has both technological and institutional determinants. Policies can therefore be designed to reap the rewards that these technological advances should be able to deliver, in terms of a more efficient information ecosystem, which presumably can be characterized by lower costs of production and dissemination and a more informed information ecosystem. Such a system is likely to be characterized by:

Policies can therefore be designed to reap the rewards that these technological advances should be able to deliver, in terms of a more efficient information ecosystem, which presumably can be characterized by lower costs of production and dissemination and a more informed information ecosystem.

- Balanced intermediation compensation, so that platforms and AI systems share value fairly with those who create information;
- Increased screening capacity, through transparency, verification infrastructure, and education; and
- A higher cost of lies, via accountability and liability regimes that make polluting the information ecosystem with lies unprofitable, or at least less profitable.

These mechanisms transform a complex and contentious topic into a set of actionable levers. The next section examines specific policy recommendations and explores whether they can ensure that AI's extraordinary capabilities support, rather than erode, the production and dissemination of truthful information.

III. Policy Recommendations

Policy reforms are naturally divided into those that increase the supply of (good) information and those that reduce the supply of mis- and disinformation. None of these are obtrusive; for the

system as a whole, they have the effect of encouraging good (welfare-enhancing) innovations, while discouraging welfare-deceasing activities, with the result that societal welfare is overall higher.

Increasing the supply of (good) information

The most important reform is ensuring that AI and platforms pay for the use of legacy media information. But simply increasing copyright protection (and eliminating the fair use exception, which AI companies have claimed gives them the right to appropriate without compensation whatever information they want from producers) won't suffice. One needs to take into account asymmetries of bargaining power—a problem addressed by Australia's bargaining code.[7] Indeed, recent analyses of "fair" allocation of platform profits[8] show that the amounts now being given by Google, for example, to legacy media are but a fraction of what a fair allocation would be.

One needs to take into account asymmetries of bargaining power—a problem addressed by Australia's bargaining code. Indeed, recent analyses of "fair" allocation of platform profits show that the amounts now being given by Google, for example, to legacy media are but a fraction of what a fair allocation would be.

Because the calculations rely on hard-to-get data, an alternative is to use simple rules, like those used in music, with proceeds distributed to originators of information on criteria reflecting both the quality and quantity of information used or produced. One of the advantages of Australia's bargaining code is that it elides that issue, leaving it to the process of bargaining between the technology or AI companies and the producers of information.

Perhaps even simpler is digital taxation, especially directed at technology platforms that transmit news and AI companies that ingest news, an increasingly popular form of taxation because of its low transactions (compliance) costs, with some of revenues devoted to

legacy media. There are multiple advantages of this digital taxation, and many governments have already adopted them.

A second related way of increasing the quality of the information ecosystem is public support for the media, recognizing that it is, as we have repeatedly noted, a public good. Crucial in the success of this intervention is that there be good, independent governance; several well-functioning democracies have (so far) shown that this is possible. Governance is crucial, because of the sensitivity of the process of allocating funds.

Media diversity is important for ensuring the availability of a range of ideas, and media concentration has become a problem in many countries, especially at the local level. Some Scandinavian countries have effectively used funds to maintain a modicum of diversity at the local level. Many governments have supported the media indirectly, through spending on advertising (e.g., posting jobs and procurement). Some recent research has questioned the efficacy of advertising, because of the difficulty in preventing political capture.

Media diversity is important for ensuring the availability of a range of ideas, and media concentration has become a problem in many countries, especially at the local level.

Public support is needed, but public production of "news" may also be desirable. Many governments do this (BBC, Canadian Broadcast Company, etc.). We've noted that the business model of the tech platforms and AI is not well aligned with societal objectives. In standard economics, it is argued that the "pursuit of self-interest leads to societal well-being." While Adam Smith identified a *force* going in that direction, he also discussed at length forces going the other way—self-interest led firms to collude, whether to raise prices or lower wages, and this went against societal well-being—Milton Friedman seems to have read

and understood part of Smith's argument, elevating the idea that firms should maximize their share value to a moral principle, now enshrined in corporate governance statutes in many jurisdictions. Yet research over the past 50 years has shown that in general, maximizing shareholder value did not maximize societal well-being; this is especially so in the markets we are examining here, where asymmetries and imperfections of information are essential. We noted that social media can increase profits by engagement through enragement—with no social value produced by the increased engagement (as it may entail engaging with mis- and disinformation) and enormous societal harms arising from the enragement (polarization, incitement to riot and to racial hatred, etc.). With good governance (and many governments have shown this possible), state provision provides a reliable "truthful" source of information, a benchmark for comparison, the effect of which may be to pull up the quality of other providers, i.e., increase the tendency to produce truthful news. Further, there is evidence that such providers increase trust in news providers more generally (there are trust spillovers to other media), and the overall demand for news increases, improving thereby the information ecosystem both in quantity and quality of the information embedded.

With good governance (and many governments have shown this possible), state provision provides a reliable "truthful" source of information, a benchmark for comparison, the effect of which may be to pull up the quality of other providers, i.e., increase the tendency to produce truthful news.

Reducing the supply and dissemination of mis- and disinformation

Here, as we already noted, the crucial problem is lack of accountability. In the US, legislation in Section 230 of the Communications Decency Act of 1996[9] essentially freed digital intermediaries from liability of what information they conveyed—giving them strongly preferential treatment

compared to standard media. This has led them to irresponsible behavior—a refusal to moderate their content, to take down scams, to prevent mis- and disinformation from going viral, and indeed employing algorithms that can be exploited to make such information go viral. During the pandemic, they showed that they had the ability to engage in content moderation, as they curbed the dissemination of dangerous anti-vax mis- and disinformation. At one time, there was hope that they would regulate themselves, but just as in banking self-regulation was a chimera, so too here.

The first reform is the repeal of Section 230 (and similar legislation elsewhere), to restore a modicum of accountability. The exception from liability was passed supposedly to support a nascent industry; that argument is hardly applicable today.

This does not suffice. There needs to be strong regulation, e.g., content moderation, and a compulsory takedown of scams and clearly false information, with punishment including loss of TV licenses and financial sanctions. Europe has led the way with its Digital Services Act (DSA),[10] but stronger regulation is still needed, even in Europe.

There needs to be strong regulation, e.g., content moderation, and a compulsory takedown of scams and clearly false information, with punishment including loss of TV licenses and financial sanctions.

There also needs to be a disclosure of algorithms used to determine what information is promoted, so they can be better assessed, for instance, on whether they are acting in a discriminatory way. Algorithms do what editors used to do: Editors decide what stories get prominence, which stories get "buried" on page 16. Algorithms do much the same (though one might say they do it on steroids.) "Promoted" articles will be widely seen. But the algorithms are especially powerful "editors," because they can target different consumers differently.

Critics of regulations like the DSA accuse it of censorship, a violation of the foundational principle of free speech, embodied in the US in the

First Amendment. But the First Amendment has never been absolute: One cannot cry fire in a crowded theater, publish pornography, or engage in slander. We've always balanced free speech with the societal harms that can be generated, with a heavy hand tilted toward the presumption of free speech. AI and the technology platforms have changed the information system, necessitating changes how we balance things *to enhance societal overall welfare.* A society with a dysfunctional information ecosystem cannot thrive. The above is the minimum that has to be done to ensure that that is the case.

Reforming ownership, preventing capture, promoting diversity

Reforms of ownership to prevent capture and promote diversity could increase both the quantity and quality of information. The standard of excessive concentration should be different from that of the usual antitrust enforcement, which focuses on market power in advertising. It is market power in the marketplace of ideas that matters. This perspective suggests more caution about joint ownership of newspapers, TV, and radio, and puts greater emphasis on diversity requirements. Public ownership, of at least part of the media space, at least in those countries where there is good governance, would eliminate some of the perverse incentives for the production of mis- and disinformation now evident.

The standard of excessive concentration should be different from that of the usual antitrust enforcement, which focuses on market power in advertising. It is market power in the marketplace of ideas that matters.

IV. TAI and Looking Ahead

The arrival of TAI promises unprecedented capabilities in reasoning, synthesis, and automation. It may one day function not only as a processor of existing information but as a generator

of new knowledge—able to conduct experiments, verify claims, and iterate hypotheses at a speed that surpasses human institutions. Yet if the lessons from our framework hold, the risks to the information ecosystem will not begin with TAI's arrival—they will accumulate on the path toward it.

If the lessons from our framework hold, the risks to the information ecosystem will not begin with TAI's arrival—they will accumulate on the path toward it.

The danger is not that TAI will suddenly destroy the market for truthful information, but that the transition itself will quietly undermine it. Long before AI systems achieve fully autonomous discovery, today's intermediaries are already absorbing audience attention, weakening attribution, and diverting revenues from original producers. As generative models become "good enough" at producing plausible content, consumers rely less on verified sources. If these trends persist, the productive base of truth creation—newsrooms, research institutions, fact-checking organizations, independent experts—may shrink beyond recovery.

This transition trap creates a policy paradox. The technologies that hold the promise of vastly increasing human knowledge may, if left unchecked, erode the very ecosystem of incentives and institutions required to realize that promise. Societies could enter the TAI era with degraded knowledge stocks, weaker public trust, and limited capacity for verification. At that point, even the most powerful systems would be drawing on a shallower pool of truthful information.

Our framework helps imagine two stylized futures, each illustrating a different equilibrium between technological capacity and institutional adaptation.

1. **AI as universal synthesizer**. AI becomes the dominant interface for information retrieval, providing near-perfect summaries of existing knowledge but creating little new fact. In absence of regulation, the main risk is a collapse of discovery incentives: Why fund original reporting or research if the value is captured elsewhere?
2. **AI as fully transformative investigator**. In this hypothetical future, TAI can test, verify, and refine its own knowledge base across disciplines. The risk here is concentration of power. A handful of actors could control the models that generate, validate, and disseminate knowledge, creating vulnerabilities to bias, capture, and systemic error. There is an especial concern about this concentration of power when the objectives of the enterprises and their owners is not well aligned with that of society as a whole; and there is a strong presumption that that is the case. Democracy requires pluralism, and maintaining pluralism will require separation of functions—between discovery, distribution, and audit—and mandatory transparency of processes through independent oversight.

The task is to align our rules with the continued production of quality information. If there is as a result a slowing down of the pace of innovation, it is a cost worth paying: The benefits of the faster arrival of TAI are far outweighed by the dangers of an ever deteriorating information ecosystem.

The same four channels remain decisive—efficiency, intermediation, screening, and the cost of lies—only at higher stakes and faster speeds. The decisions societies make now—on intellectual property, attribution, liability, and public support—will determine whether future technologies expand our collective knowledge or erode it. The task is to align our rules with the continued production of quality information. If there is as a result a slowing down of the pace of innovation, it is a cost worth paying: The benefits of the faster arrival of TAI are far outweighed by the dangers of an ever deteriorating information ecosystem.

A resilient, high-quality, plural, and transparent information ecosystem is the foundation on which both democracy and technological progress ultimately depend.

1. John H. Cochrane, "AI, Society, and Democracy: Just Relax," in Erik Brynjolfsson, Alex Pentland, Nathaniel Persily, Condoleezza Rice, and Angela Aristidou, eds., *The Digitalist Papers* (Stanford Digital Economy Lab, 2024): 127–141, https://www.digitalistpapers.com/essays/ai-society-and-democracy-just-relax.
2. Joseph E. Stiglitz and Màxim Ventura-Bolet, "The Impact of AI and Digital Platforms on the Information Ecosystem," Working Paper No. 34318 (National Bureau of Economic Research, October 2025), https://doi.org/10.3386/w34318.
3. Kenneth J. Arrow, "Economic Welfare and the Allocation of Resources for Invention," in Richard Nelson, ed., *The Rate and Direction of Inventive Activity: Economic and Social Factors* (Princeton University Press, 1962), https://www.jstor.org/stable/j.ctt183pshc; Joseph E. Stiglitz, "Information and Economic Analysis," in Michael Parkin and A.R. Nobay, eds., *Current Economic Problems* (Cambridge University Press, 1975): 27–52; Joseph E. Stiglitz, "The Theory of Screening, Education and the Distribution of Income," *American Economic Review* 65, no. 3 (1975): 283–300, https://doi.org/10.7916/D8PG22PM; Joseph E. Stiglitz, "On the Microeconomics of Technical Progress," in Jorge M. Katz, ed., *Technology Generation in Latin American Manufacturing Industries* (St. Martin's Press, 1987): 56–77, https://doi.org/10.1007/978-1-349-07210-1_3; Joseph E. Stiglitz, "Knowledge as a Global Public Good," in Inge Kaul, Isabelle Grunberg, and Marc A. Stern, eds., *Global Public Goods: International Cooperation in the 21st Century* (Oxford University Press, 1999): 308–25, https://doi.org/10.1093/0195130529.003.0015; Joseph E. Stiglitz, "The Media: Information as a Public Good," paper presented to the Pontifical Academy of Social Sciences, May 10, 2021, https://business.columbia.edu/sites/default/files-efs/imce-uploads/Joseph_Stiglitz/The%20Media%20Informatino%20as%20a%20Public%20Good%20Slides.pdf.
4. Although the overall effect is evident—many information producers have shut down or downsized in recent years—this effect varies across types of information and producers (see Lyu et al., 2025) and can depend on the relative strength of two opposing forces (as found in Jeon and Nasr, 2016). There is a substitution effect, whereby information aggregators divert traffic away from original news sources, and an expansion effect, whereby these aggregators enhance overall demand for information, potentially increasing exposure and visits to original producers (as found in Calzada and Gil 2020, and Athey et al. 2021). Currently, effective referrals from AI to the original sources from which the information is derived appear much weaker for AI than for the technology platforms, suggesting that the adverse effects of AI may be much greater (see Chapekis and Lieb 2025). Liang Lyu, James Siderius, Hannah Li, Daron Acemoglu, Daniel Huttenlocher, and Asuman Ozdaglar, "Wikipedia Contributions in the Wake of ChatGPT," in *WWW '25: Companion Proceedings of the ACM on Web Conference 2025* (ACM, May 2025): 1176–1179, https://doi.org/10.1145/3701716.3715543; Doh-Shin Jeon and Nikrooz Nasr, "News Aggregators and Competition Among Newspapers on the Internet," *American Economic Journal: Microeconomics* 8, no. 4 (2016): 91–114, https://doi.org/10.1257/mic.20140151; Joan Calzada and Ricard Gil, "What Do News Aggregators Do? Evidence from Google News in Spain and Germany," *Marketing Science* 39, no. 1 (January 2020): 134–67, https://doi.org/10.1287/mksc.2019.1150;

Susan Athey, Markus Möbius, and Jeno Pal, "The Impact of Aggregators on Internet News Consumption," Working Paper No. 28746 (National Bureau of Economic Research, April 2021), https://doi.org/10.3386/w28746; Athena Chapekis and Anna Lieb, "Google Users Are Less Likely to Click on Links when an AI Summary Appears in the Results," Pew Research Center, July 22, 2025, https://www.pewresearch.org/short-reads/2025/07/22/google-users-are-less-likely-to-click-on-links-when-an-ai-summary-appears-in-the-results/.

5. There are two further parameters not explored in this paper but which we hope to take up in future research: the effectiveness of reputation mechanisms and the extent of targeting.
6. Sanford Grossman and Joseph E. Stiglitz, "On the Impossibility of Informationally Efficient Markets," *American Economic Review* 70, no. 3 (1980): 393–408, https://www.jstor.org/stable/1805228.
7. Treasury Laws Amendment (News Media and Digital Platforms Mandatory Bargaining Code) Act 2021, No. 21, 2021, https://www.legislation.gov.au/C2021A00021/asmade/text.
8. Using standard methodologies, e.g., Mateen et al. 2023 (based on principles put forward by Nalebuff 2021). Haaris Mateen, Haris Tabakovic, Patrick Holder, and Anya Schiffrin, "Paying for News: What Google and Meta Owe U.S. Publishers," working paper (Initiative for Policy Dialogue, Columbia University, November 2023), https://ipdcolumbia.org/wp-content/uploads/2024/08/Paying-for-News-What-Google-and-Meta-Owe-US-Publishers-—-Draft-Working-Paper.pdf; Barry Nalebuff, "A Perspective-Invariant Approach to Nash Bargaining," *Management Science* 67, no. 1 (2021): 577–593, https://doi.org/10.1287/mnsc.2019.3547.
9. Communications Decency Act (CDA), 47 U.S.C. § 230.
10. Regulation (EU) 2022/2065 of the European Parliament and of the Council of 19 October 2022 on a Single Market For Digital Services and amending Directive 2000/31/EC (Digital Services Act) O.J. (L 277).