UNITED STATES DISTRICT COURT
SOUTHERN DISTRICT OF NEW YORK

IN RE: OpenAI, Inc., COPYRIGHT INFRINGEMENT LITIGATION This Document Relates To: 1:23-cv-11195-SHS-OTW 1:24-cv-01515-SHS-OTW 1:24-cv-03285-SHS-OTW 1:24-cv-04872-SHS-OTW 1:25-cv-04315-SHS-OTW	25-md-3143 (SHS) (OTW)

[PROPOSED] BRIEF OF PROFESSOR JOSEPH STIGLITZ AS AMICUS CURIAE IN SUPPORT OF PLAINTIFFS

STATEMENT REGARDING AUTHORSHIP AND FUNDING

No party's counsel authored this brief in whole or in part. No party, party's counsel, or person other than Professor Joseph Stiglitz and his counsel contributed money intended to fund the preparation or submission of this brief.

TABLE OF CONTENTS

I. INTEREST OF AMICUS CURIAE 1
II. INTRODUCTION 3
III. ARGUMENT 5
A. Framework for Considering AI's Effects on the Information Ecosystem 6
1. The Pre-AI Information Ecosystem 6
2. Motivation and Effects of AI 7
3. The Model and Equilibrium Outcomes 8
4. Risks When Parameters Move Toward Extremes 9
B. In the Budding AI Era, Courts Should Apply the Fair Use Doctrine Sparingly 11
1. The Fourth Factor Weighs Decisively Against Fair Use Because the Content Generated by Defendants' Use Substitutes for Plaintiffs' Works 12
2. The First Factor Weighs Against Fair Use Because Defendants Employ Plaintiffs' Works for the Same Commercial Purpose 13
3. The Asserted Public Benefits do not Justify the Use, Because the Use's Net Effect Degrades the Public's Access to Reliable Information 14
C. Fairness Considerations 15
IV. CONCLUSION 16

TABLE OF AUTHORITIES

Cases

Andy Warhol Found. for the Visual Arts, Inc. v. Goldsmith
598 U.S. 508 (2023) 13, 14

Associated Press v. Meltwater U.S. Holdings, Inc.
931 F. Supp. 2d 537 (S.D.N.Y. 2013) 14

Authors Guild v. Google, Inc.
804 F.3d 202 (2d Cir. 2015) 11, 12

Authors Guild, Inc. v. HathiTrust
755 F.3d 87 (2d Cir. 2014) 11

Campbell v. Acuff-Rose Music, Inc.
510 U.S. 569 (1994) 12

Fox News Network, LLC v. TVEyes, Inc.
883 F.3d 169 (2d Cir. 2018) 12

Harper & Row, Publishers, Inc. v. Nation Enters.
471 U.S. 539 (1985) 11, 12

Sony Corp. of Am. v. Universal City Studios, Inc.
464 U.S. 417 (1984) 11

Thomson Reuters Enter. Ctr. GmbH v. Ross Intelligence Inc.
No. 25-2153, 2026 WL 2932557 (3d Cir., Sept. 29, 2026) 13

Statutes

17 U.S.C. § 107 12

U.S. Const. art. I, § 8, cl. 8 15

Secondary Sources

Barry Nalebuff, *Split the Pie: A Radical New Way to Negotiate* (2022) 16

John H. Cochrane, *AI, Society, and Democracy: Just Relax* in *The Digitalist Papers: Artificial Intelligence and Democracy in America* (Erik Brynjolfsson et al. eds., 2024) 5

Joseph E. Stiglitz & Maxim Ventura-Bolet, *Information in the Age of AI: Challenges and Solutions* in *The Digitalist Papers, Vol. 2: The Economics of Transformative AI* (Stanford Digital Economy Lab 2025) 6

Joseph E. Stiglitz & Maxim Ventura-Bolet, *The Impact of AI and Digital Platforms on the Information Ecosystem* (Nat'l Bureau of Econ. Research, Working Paper No. 34318, 2025) .. 6, 13, 14, 15

Sanford J. Grossman & Joseph E. Stiglitz, *On the Impossibility of Informationally Efficient Markets*, 70 Am. Econ. Rev. 393 (1980) 10

I. INTEREST OF AMICUS CURIAE

Professor Joseph Stiglitz ("Stiglitz") is a University Professor at Columbia University with joint appointments in the Faculty of Arts and Sciences (Department of Economics), the Graduate School of Business, and the School of International and Public Affairs.[1] Prior to assuming this position, Stiglitz held professorships at Stanford University, Yale University, Princeton University, the Massachusetts Institute of Technology, and the University of Oxford, where he taught a wide variety of graduate and undergraduate courses in economics and finance.

Stiglitz has received numerous fellowships and honors over his career. In 2001, Stiglitz was awarded the Nobel Memorial Prize in Economics for his work on Information Economics. This work includes the study of how information asymmetries affect economic behavior, the determination of the conditions under which efficient sharing of risk occurs, the economics of financial markets, and the production of knowledge. His particular concern in this brief is how alternative legal frameworks, including the application of the fair use doctrine, affect the information ecosystem. He also has a particular lifelong concern on issues of equity. This brief reflects some of his work and that of his colleagues on the equity implications of legal frameworks.

Over the course of his career, Stiglitz has published hundreds of peer-reviewed articles, written or edited more than 50 academic and popular books, testified several times before Congress, and written numerous opinion pieces for newspapers and magazines. His publications and research have extended into many different areas, including macroeconomics and monetary theory, development economics and trade theory, public and corporate finance, industrial organization and rural organization, welfare economics, and income and wealth distribution.

From 1993 to 1997, Stiglitz served as a member of President Clinton's Council of Economic Advisers, and from 1995 to 1997, as Chairman of the Council and a member of the President's Cabinet. As Chairman and Cabinet Member, Stiglitz was heavily involved in formulating fiscal policy, sustainable economic policies, financial sector regulation and banking

[1] Stiglitz's CV and biography can be found at https://business.columbia.edu/stiglitz/about.

policy, and coordinating policy with the U.S. Treasury. Within the Council, he was responsible for formulating views on intellectual property, especially relevant at the time he served because of the Trade-Related Aspects of Intellectual Property Rights ("TRIPS") Agreement, administered by the WTO, which was entered during his tenure. He was also the responsible person working with the Office of Science and Technology Policy ("OSTP") on science and technology policy and assessing the economic and societal impacts of different policies.

From 1997 to 2000 he served as senior vice president and chief economist of the World Bank, and in that position was responsible for advising multiple countries on various aspects of science and technology policy, including on intellectual property matters.

Professor Stiglitz's unique scientific credentials are reflected in his being one of the few economists elected to the Royal Society of London for Improving Natural Knowledge, the oldest continuously existing scientific academy in the world. He has also been elected as a fellow or member of the National Academy of Sciences, the American Academy of Arts and Science, the Econometric Society, the British Academy, and academies of sciences in other countries, including a lifelong member of the Papal Academy of Social Sciences. He has received more than 60 honorary degrees. He was elevated to the rank of *officier* in the French Legion of Honor in 2012.

Stiglitz has provided expert testimony in various fora throughout the United States and before foreign courts and international tribunals, and has submitted amicus curiae briefs before multiple courts. He has published and lectured widely on intellectual property, including in economics, law, and medical journals.

Stiglitz has served or is serving currently on many commissions and advisory committees addressing a myriad of economic policy issues, both in the United States and abroad. Relevant to this case, Stiglitz serves on the Advisory Board of the International Association for Safe & Ethical AI, whose leadership includes many of those responsible for the creation of AI. Stiglitz's analysis of the effects of AI on our economic system began in the early days of AI, including a paper on the distributive effects of AI in 2017 (subsequently published in 2019), and a lecture on AI at the

Royal Society in 2018, accompanied by meetings with leading AI researchers, including from Deep Mind.

II. 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 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 in which 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? Is our information ecosystem improving? *It does not follow from the fact that technology is improving that the information ecosystem is improving. For that to happen there must be appropriate laws and regulations guiding technology. That is a critical issue in this case.*

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 exacerbate them.

There are three separate problems, each related to digital platforms, but even more so to AI.

The first problem is an undersupply of good information. Even with strong intellectual property laws, there are knowledge spillovers (externalities). This is true of the kind of information that we think of as "news" as well as 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 available knowledge. Others take 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. Accordingly, there is no presumption of optimality in the production, transmission and usage of information.

Technology platforms and AI have made matters worse, because they are not just competitors of the traditional producers and disseminators of information, but they are, in part, competitors who take information from these traditional producers of information and 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 pollute the information ecosystem don't pay for the consequences, the damage that occurs as a result of a poorer information ecosystem, or the resources that have to be deployed to separate 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 their users up. 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 pre-existing 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 are an ample 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 law and 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.[2]

That view is totally misguided. It simply assumes that all innovation is welfare enhancing, when it should be obvious that is not the case: innovations in creating and exploiting market power are welfare decreasing. There is a real risk that these technological innovations too may be welfare decreasing in the form of harming the information ecosystem. In other words, technological advances that have the power to improve the information ecosystem may have just the opposite effect.

However, with the right legal and regulatory framework, innovation in these areas may be beneficial to society as a whole, as opposed to just increasing the profits of the tech companies. Further, steering innovation in ways that enhance the information ecosystem (rather than worsen it) may enhance the overall pace of productive innovation. Moreover, the rules that would lead to a better ecosystem are consistent with a *fairer* sharing of the gains from the advances in technology.

III. ARGUMENT

The analysis below is based upon and describes a simplified version of a model Stiglitz and Maxim Ventura-Bolet developed and described in a 2025 National Bureau of Economic Research working paper. *See* Joseph E. Stiglitz & Maxim Ventura-Bolet*, The Impact of AI and Digital Platforms on the Information Ecosystem* (Nat'l Bureau of Econ. Research, Working Paper

[2] *See, e,g.*, John H. Cochrane, *AI, Society, and Democracy: Just Relax*, in *The Digitalist Papers: Artificial Intelligence and Democracy in America* (Erik Brynjolfsson et al. eds., 2024).

No. 34318, 2025) ("Stiglitz & Ventura-Bolet").[3] Stiglitz appends that paper to this brief.[4] In short, this analysis shows that when AI intermediaries absorb the attention and revenue that once flowed to producers of reliable information, the incentive to produce that reliable information erodes, and the stock of accurate information available to society may shrink over time, even if the technology of processing and transmitting information improves. As the old expression puts it, "garbage in, garbage out." This dynamic must be considered in the Court's fair use analysis. *See* Section III.B, *infra*.

A. Framework for Considering AI's Effects on the Information Ecosystem

The analytical framework described below distinguishes the channels through which AI enhances the production and dissemination of high-quality information from those through which it undermines them.

1. 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. (Economists say such 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.

---

[3] This paper was just published in *The Economics of Transformative AI* (eds. Ajay Agrawal, Erik Brunjolfsson and Anton Korinek, University of Chicago Press).

[4] Stiglitz and Bentura-Bolet published a short article building upon the NBER working paper, from which this brief borrows. *See* Joseph E. Stiglitz & Maxim Ventura-Bolet, *Information in the Age of AI: Challenges and Solutions*, in *The Digitalist Papers, Vol. 2: The Economics of Transformative AI* (Stanford Digital Economy Lab 2025).

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 with incentives for its creation and dissemination. And for much of the twentieth 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.

2. Motivation and Effects of AI

How will AI affect the information ecosystem? AI introduces several powerful and interrelated shocks:

Improved production and dissemination. AI systems dramatically enhance the efficiency of processing and transmitting information, allowing content to reach broader audiences (and, in some cases, more targeted audiences) faster and at a lower cost than before.

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.[5]

Falling relative cost of lies. Generative models make the creation of persuasive misinformation inexpensive compared with producing verified truth.

[5] This dynamic can produce a substitution effect, whereby information aggregators divert traffic away from original news sources, and an expansion effect, whereby aggregators enhance overall demand for information, potentially increasing exposure and visits to original producers. But effective referrals to original sources appear weaker for AI than for other technology platforms. *See, e.g.*, News Plaintiffs' Combined Summary Judgment Brief (Redacted) (Dkt. No. 2108) at 68-69 (discussing significantly reduced click-through rates for Bing Chat as compared to Bing Search). This suggests that the adverse effects of AI are much greater than those of earlier technologies.

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, with evidence and theory that, at least in the short run, deception may dominate.

3. The Model and Equilibrium Outcomes

In modelling these dynamics, one may assume there are two types of agents—producers and consumers. Their interactions determine the 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, *i.e.*, where incentives to produce at least certain kinds of socially desirable information are so weak that such information is not produced at all.

As with any model, this 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. Of course, in the real-world content can fall on a continuum between perfect "truth" and perfect "untruth." The results of the analysis can be shown to be robust to these simplifying assumptions.

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").

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.

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.

4. Risks When Parameters Move Toward Extremes

Risks emerge when the parameters described above move toward their extremes. (It is important to emphasize, however, that the problems identified here do not arise just as the extremes.) Each has intuitive implications.

First, when intermediation (obtaining news through technology platforms or AI) is 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. *See* Sanford J. Grossman & Joseph E. Stiglitz, *On the Impossibility of Informationally Efficient Markets*, 70 Am. Econ. Rev. 393 (1980). 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. If the AI systems become anywhere near as good as their advocates claim and the market seems to believe, the risk (under current practice) would be a collapse of the quality of the information ecosystem.

Next, 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.

Finally, low screening capacity means consumers cannot reliably distinguish true from false information. Demand for quality weakens, and misinformation becomes self-reinforcing: as trust deteriorates (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. While AI may enable the better detection of fakes and false information, the ability to deceive (of which there are a large number of ways) may exceed the ability to detect false information, a situation in which we seem currently to be.

B. In the Budding AI Era, Courts Should Apply the Fair Use Doctrine Sparingly

The framework and analysis set forth above has multiple policy implications. For example, recognizing that a high-quality information ecosystem is a public good might suggest policy changes designed to increase the supply of good information and policies intended to reduce the supply and dissemination of mis- and disinformation. But those broad policy questions are not before this Court. Instead, Stiglitz focuses on the impact a too broad reading of the fair use doctrine could have on the information ecosystem in the AI era.

Fair use is an affirmative defense that "allows the public to draw upon copyrighted material without the permission of the copyright holder in certain circumstances" when "necessary to fulfill copyright's very purpose, to promote the Progress of Science and useful Arts." *Authors Guild, Inc. v. HathiTrust*, 755 F.3d 87, 95 (2d Cir. 2014) (cleaned up). Put another way, copyright "is intended to motivate the creative activity of authors" by securing them a return on their labor, *Sony Corp. of Am. v. Universal City Studios, Inc.*, 464 U.S. 417, 429 (1984), and the fair use inquiry asks whether a given use serves or undermines that goal. *See Authors Guild v. Google, Inc.*, 804 F.3d 202, 212 (2d Cir. 2015); *Harper & Row, Publrs. v. Nation Enters.*, 471 U.S. 539, 558 (1985). The analysis described above shows that Defendants' use subverts that goal.

The four statutory factors to be considered when determining whether the fair use doctrine applies include:

(1) the purpose and character of the use, including whether such use is of a commercial nature or is for nonprofit educational purposes;

(2) the nature of the copyrighted work;

(3) the amount and substantiality of the portion used in relation to the copyrighted work as a whole; and

(4) the effect of the use upon the potential market for or value of the copyrighted work.

17 U.S.C. § 107. The litigants in this case have discussed these factors in great depth, and other *amici* are expected to do the same. Stiglitz focuses on the fourth factor, *i.e.*, the effect of the use on the potential market for or value of the copyrighted work, and to a lesser extent the first factor, *i.e.*, whether the purportedly fair use has a further purpose or different character.

1. The Fourth Factor Weighs Decisively Against Fair Use Because the Content Generated by Defendants' Use Substitutes for Plaintiffs' Works

The fourth factor is "undoubtedly the single most important element of fair use." *Harper & Row*, 471 U.S. at 566. A court considering this factor asks not only whether the challenged use harms the market for the original, but also "whether unrestricted and widespread conduct of the sort engaged in by the defendant" would "result in a substantially adverse impact on the potential market." *Campbell v. Acuff-Rose Music, Inc.*, 510 U.S. 569, 590 (1994). In the Second Circuit, the question is whether the secondary use offers a competing substitute that deprives the rightsholder of significant revenue. *Authors Guild*, 804 F.3d at 223; *see also Fox News Network, L.L.C. v. TVEyes, Inc.*, 883 F.3d 169, 179-80 (2d Cir. 2018) (fourth factor favored rightsholder where service gave users Fox's content in a form that displaced Fox's own revenue and licensing market).

The model described above explains why this use of copyright-protected material gives rise to substitution that is not just a marginal harm, but a structural one. Producers of reliable information fund their work from the attention of readers (through subscription fees or advertising revenue) who visit them directly. When an intermediary captures a larger share of that attention,

the model shows that the incentive to produce new information falls and the stock of accurate information declines with it. *See* Stiglitz & Ventura-Bolet at 14 & fig. 4. Generative AI aggravates this problem in ways that earlier aggregators did not. Search engines and social media platforms typically attributed their sources and arguably sent some traffic to the content owners (though in the case of social media, the traffic sent to content owners was far smaller than the traffic diverted away). But the diversion is even greater for an AI system that synthesizes many sources into a single answer. Users appear to "click through" at rates lower than with earlier technologies. *See, e.g.*, News Plaintiffs' Combined Summary Judgment Brief (Redacted) (Dkt. No. 2108) at 68-69. This is the textbook definition of a market substitute.

If this practice becomes the norm, it could push the information ecosystem past a tipping point to "information collapse," in which reliable information would not be produced because the producers could not recoup the cost. Stiglitz & Ventura-Bolet at 9-10, 14.

2. The First Factor Weighs Against Fair Use Because Defendants Employ Plaintiffs' Works for the Same Commercial Purpose

Under *Andy Warhol Foundation for the Visual Arts, Inc. v. Goldsmith*, 598 U.S. 508 (2023), the first factor asks whether the use has a further purpose or different character, which is a question of degree that must be weighed against the use's commercial nature. *Id.* at 528-33. When the original and the copy serve substantially the same purpose and the copy is commercial, the first factor likely weighs against fair use absent some other justification. *Id.* at 532-33.

Here Defendants' use shares Plaintiffs' purpose. The News Plaintiffs' works exist to inform the public. Defendants' products answer users' questions by delivering the same information, drawn from the same works, in a commercial product that competes for the same readers. *See* News Plaintiffs' Combined Summary Judgment Brief (Redacted) (Dkt. No. 2108) at 52-55 (discussing evidence, much from the AI companies themselves, that Defendants' products' content substitutes for news). Just last month the Third Circuit held that copying a publisher's content to build a competing AI research tool was minimally transformative at best, because the AI product served the same purpose as the publisher's product. *Thomson Reuters Enter. Ctr. GmbH v. Ross Intel. Inc.*, No. 25-2153, 2026 WL 2932557, at *6-7 (3d Cir., Sep. 29, 2026).

Authority from the Second Circuit and the Southern District of New York is in accord. *See Hachette Book Grp., Inc. v. Internet Archive*, 115 F.4th 163, 181 (2d Cir. 2024) (delivering copyrighted works in a new, more convenient form for the same purpose is not transformative); *AP v. Meltwater U.S. Holdings, Inc.*, 931 F. Supp. 2d 537, 556-57 (S.D.N.Y. 2013) (news monitoring service that republished excerpts of AP articles was not transformative).

The efficiency of Defendants' technology does not provide the "further purpose" *Warhol* requires. As described above and in Stiglitz & Ventura-Bolet, the two relevant effects of AI here are (1) a potential gain in the efficiency by which existing information is transmitted and processed, possibly combined with information and dis- and misinformation from others, and (2) a loss in the incentive to produce new information. Stiglitz & Ventura-Bolet at 11-13. The model shows that the first effect can be swamped by the second.[6] While the public may initially perceive a gain, as producers of news exit, the stock of high-quality information shrinks and the overall pool of accurate knowledge may end up shallower than before. A use whose efficiency gains come at the cost of the disappearance of the underlying works is not transformative. It is just a more efficient way of taking.

3. The Asserted Public Benefits do not Justify the Use, Because the Use's Net Effect Degrades the Public's Access to Reliable Information

Defendants have touted the public benefits of their technology. But as the Second Circuit recognized on a similar argument, that argument is "shortsighted:"

> Within the framework of the Copyright Act, [Internet Archive's] argument regarding the public interest is shortsighted. True, libraries and consumers may reap some short-term benefits from access to free digital books, but what are the long-term consequences? If authors and creators knew that their original works could be copied and disseminated for free, there would be little motivation to produce new works. And a dearth of creative activity would undoubtedly negatively impact the public. This is the reality that the Copyright Act seeks to avoid.

Hachette Book Group, 115 F.4th at 195.

[6] Even in conducting the first, there may be a loss of efficiency, *e.g.*, arising from mis- and disinformation, but one of which the user is not fully aware, so that substitution against the original producer still occurs.

Information is a public good. Once published, it spreads to people who never paid for it, so it will be undersupplied unless the law secures its producers a return. Stiglitz & Ventura-Bolet at 1-2. Copyright is one of the few legal mechanisms that permits information producers to be compensated for their efforts to generate it. Allowing AI systems to sidestep copyright protections does not make good information more abundant. Rather, it will likely make reliable information scarcer.

As Stiglitz and Ventura-Bolet have shown, generative AI dramatically lowers the cost of producing false or low-quality content relative to the cost of producing accurate reporting, and may impair information consumers' ability to distinguish between the two. *Id.* at 17-18. As the model described above shows, these shifts can increase the share of misinformation in the information ecosystem and lower the relative stock of truthful information, together driving the ecosystem toward a state in which little or no truthful information is produced. *Id.* at 16-18. If producers of original content (especially of high-quality content entailing investments in gathering information (investigative reporting), writing, editing, and fact checking) cannot protect the fruit of their labor from uncompensated use, AI systems will increasingly pull from a degraded pool of sources. The "public benefit" Defendants offer is therefore self-undermining, as their products draw their value from the very works whose production their uncompensated use discourages. A use that depletes the pool of information from which it draws does not "promote the Progress of Science and useful Arts."[7] U.S. Const. art. I, § 8, cl. 8.

C. Fairness Considerations

While the analysis above focuses on incentives, economists often address the issue of *equitable distribution* of the benefits of an innovation in the context of interdependence. While detailed research into equitable distribution has occurred in the context of legacy media and the

[7] It is noteworthy that the United Kingdom, while having a different legal framework, came to much the same conclusion. There, the question was whether to open an exception for AI—would such an exception serve progress and the wellbeing of society? After extensive deliberation, the conclusion was no.

social media and search platforms, the principles extend in a straightforward way to AI. A hypothetical illustrates this. Assume that without the use of legacy media, a social media platform is worth $1 billion, but with appropriating the content of legacy media, it's worth $5 billion. Assume that the social media substitutes for the legacy media, so that before social media the legacy media was worth $2 billion, but after emergence of social media it's worth $1 billion. The social media platform has increased *joint value* from $3 billion to $6 billion. Ignoring that the legacy media's value has been cut in half (and the adverse effect that would have on the production of knowledge) the social media *employing legacy media* has increased total "value" by $3 billion. It is clear from this example is that the social media *needs* the legacy media (it increases their value by $4 billion). Economists have asked, "What is a fair (or equitable) division of the surplus that can only be achieved by their cooperating?" The answer provided by economists, articulated most clearly by Barry Nalebuff, is to *split equally the increase in the size of the pie*, that is, split the $3 billion equally, so the legacy media get in the end $2.5 billion and the social media get in total the other $3.5 billion. *See, e.g.*, Barry Nalebuff, *Split the Pie: A Radical New Way to Negotiate* (2022).

What the AI companies are proposing, in effect, with unchecked stealing of the product of legacy media is that the legacy media get $1 billion (the crumbs they are left with after AI garners for itself everything else), and they take the entire rest ($5 billion). The analysis above emphasizes the extreme adverse effects on incentives and therefore on the quality of the information ecosystem. But even if these effects weren't present, what the AI companies are proposing is grossly inequitable, by any standard within economics. This Court should not permit the fair use doctrine to excuse the AI companies' appropriation. Rather, content owners and AI companies should negotiate licenses to equitably split the benefits from this innovation and the content upon which it depends.

IV. CONCLUSION

AI 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, gather new information, conduct

complicated investigative reporting, and iterate hypotheses at a speed that surpasses human institutions. We are not in that world today. Yet if the lessons from the framework described above hold, the risks to the information ecosystem will accelerate as AI use becomes more and more prevalent. For the foreseeable future, until this nirvana arrives, without the appropriate legal framework, for all of its positive potential, there is a significant risk that AI will actually lead to a deterioration in the information ecosystem.

The danger is not that AI 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 and synthesizing what they have stolen from other sources, 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 AI 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.

AI has the potential to become the dominant interface for information retrieval, providing near-perfect summaries of existing knowledge but creating little new original content. The primary risk here is a collapse of discovery incentives: why fund original reporting or research if the value is captured elsewhere?

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—as relevant here, on intellectual property—will determine whether future technologies expand our collective knowledge or erode it. Our rules should be aligned with the continued production of

quality information. A resilient, high quality, plural, and transparent information ecosystem is the foundation on which both democracy and technological progress ultimately depend. This Court should bear these dynamics in mind when considering the fair use doctrine to the conduct at issue in this case.

DATED: October 30, 2026.

Respectfully submitted,

ARETE LAW GROUP PLLC

By: */s/ Jeremy Roller*
Jeremy E. Roller, WSBA No. 32021
600 University Street, Suite 2420
Seattle, WA 98101
Phone: (206) 428-3250
jroller@aretelaw.com

Attorneys for [Proposed] Amicus Curiae Professor Jospeh Stiglitz

CERTIFICATE OF COMPLIANCE WITH WORD-COUNT LIMITATION

Pursuant to Local Civil Rule 7.1(c), the undersigned certifies that this brief complies with the word-count limitation set by the Court's Order of August 7, 2026 (Dkt. No. 1663), which provides that a proposed amicus brief shall be no more than 6,000 words. As calculated by the word-count function of the word-processing program used to prepare this brief, the brief contains 5,986 words, excluding the statement regarding authorship and funding, caption, table of contents, table of authorities, signature block, and this certificate.

Dated: October 30, 2026.

/s/ Jeremy Roller
Jeremy E. Roller

Joseph E. Stiglitz & Maxim Ventura-Bolet
The Impact of AI and Digital Platforms on the Information Ecosystem
(Nat'l Bureau of Econ. Research, Working Paper No. 34318, 2025)

NBER WORKING PAPER SERIES

THE IMPACT OF AI AND DIGITAL PLATFORMS ON THE INFORMATION ECOSYSTEM

Joseph E. Stiglitz
Maxim Ventura-Bolet

Working Paper 34318
http://www.nber.org/papers/w34318

NATIONAL BUREAU OF ECONOMIC RESEARCH
1050 Massachusetts Avenue
Cambridge, MA 02138
October 2025

We are thankful to Anton Korinek, Wei Li, Andrew Souther and numerous participants at the Economics of TAI Workshop Fall 2025, Sciences Po and Columbia University for their useful comments. Ventura-Bolet gratefully acknowledges financial support from the Bank of Spain through a doctoral fellowship. The views expressed herein are those of the authors and do not necessarily reflect the views of the National Bureau of Economic Research.

NBER working papers are circulated for discussion and comment purposes. They have not been peer-reviewed or been subject to the review by the NBER Board of Directors that accompanies official NBER publications.

© 2025 by Joseph E. Stiglitz and Maxim Ventura-Bolet. All rights reserved. Short sections of text, not to exceed two paragraphs, may be quoted without explicit permission provided that full credit, including © notice, is given to the source.

The Impact of AI and Digital Platforms on the Information Ecosystem
Joseph E. Stiglitz and Maxim Ventura-Bolet
NBER Working Paper No. 34318
October 2025
JEL No. D8, D83, O33

ABSTRACT

We develop a tractable model to study how AI and digital platforms impact the information ecosystem. News producers — who create truthful or untruthful content that becomes a public good or bad —earn revenue from consumer visits. Consumers search for information and differ in their ability to distinguish truthful from untruthful information. AI and digital platforms influence the ecosystem by: improving the efficiency of processing and transmission of information, endangering the producer business model, changing the relative cost of producing misinformation and altering the ability of consumers to screen quality. We find that in the absence of adequate regulation (accountability, content moderation, and intellectual property protection) the quality of the information ecosystem may decline, both because the equilibrium quantity of truthful information declines and the share of misinformation increases; and polarization may intensify. While some of these problems are already evident with digital platforms, AI may have different, and overall more adverse, impacts.

Joseph E. Stiglitz
Columbia University
and NBER
jes322@columbia.edu

Maxim Ventura-Bolet
Columbia University
mv3010@columbia.edu

1 Introduction

AI has been heralded as opening up 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, 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.

But for a wide range of areas, AI could potentially have large negative effects on the information ecosystem. In areas of the social sciences and even in some areas of the sciences, such as evolutionary biology, the world is ever changing in ways that are hard to predict, at least at any level of AI currently or likely to be available. In these areas, it is imperative to constantly add information about the current state of the world.

Information, though, is a public good, in the sense defined by Samuelson (Arrow (1962), Stiglitz (1975, 1986, 1999, 2021)). Even with strong intellectual property laws, there are knowledge spillovers, and this is as true of the kind of information that we think of as "news" as 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 which is available. Others take out from this pool, and neither those who use the pool of knowledge (Stiglitz (2014)) nor those who contribute to it pay or receive compensation commensurate with marginal values, and accordingly there is no presumption of optimality in the production, transmission and usage of information. Indeed, the argument we have just given suggests that there will *typically* be an undersupply of information, unless there is public intervention.[1]

[1]In our complex decentralized interdependent and interrelated society/economy, information production and acquisition is widely dispersed, with virtually everyone in our economy engaged in the process in one way or the other. Inevitably there is some duplicative research. Because there are real costs of transmitting information and acquiring information, both from others and about what information others are gathering, in some instances, duplicative production of information may be efficient, and this is especially so when the output of research is not fully accurate (as is often the case). Further, much information acquisition is associated with rent seeking, for instance returns one obtains simply by having information earlier than others, or expenditures on innovation the profits from which originates mainly by acquiring some part of the (possibly monopoly) rents of others. In these cases, there may in fact be excessive expenditures on information. There is also a large literature looking at expenditures on information acquisition from this perspective. See Alchian and Demsetz (1972), Stiglitz (1975, 1989), Aghion and Howitt (1992).

There is a further problem: one has to ascertain the quality of information, and not only may there not be incentives for honesty (an essential idea within the screening/signaling literature), there can be incentives for dishonesty, unless curbed, for instance, by fraud laws.[2] 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.

1.1 AI, digital platforms, and a deterioration of the information ecosystem

Even before the advent of AI, there was a concern that the digital platforms (search engines and social media) might lead to a deterioration in the quality of the overall information ecosystem, even if all information were truthful (i.e. in the absence of mis and disinformation).

The underlying idea of our paper is simple: our economy provided a peculiar, but workable, solution to the public-good problem associated with the production of news. Newspapers and magazines produced information that attracted readers, and readers attracted advertising.[3] The hope was that by providing high quality and timely news, an outlet would attract more readers, which would attract more advertisers.[4] The uncomfortable marriage between advertising and the production and dissemination of news information continued until the arrival of the platforms, which in short order were able to garner for themselves "eyeballs," partially by appropriating information gathered by the legacy media.

The acquisition of "eye-balls" by AI or digital platforms would have, on its own, reduced incentives for legacy media to produce information. And since the social media can and do acquire eye-balls without producing news, the reduced production of new information means that, even if the digital platforms improve the efficiency of the transmission of information, it may, under quite general conditions, result in a worse information ecosystem, one where

[2]Until the advent of social media, information economics spent little time discussing the problem of mis and disinformation or the role of fraud laws. It was simply assumed that the information that individuals told was true—they told the truth, nothing but the truth, but not necessarily the whole truth. An exception is Greenwald and Stiglitz (1992).

[3]The business models typically combined subscriptions with advertising. In the model presented here, we do not deal with subscriptions. There is a broad consensus on the limitations of the subscription business model, but a discussion of this would take us beyond this paper.

[4]That hope was partially dashed by Rupert Murdoch, who first showed that more readers could be attracted by entertainment that by informative and timely news coverage; and even more so by Fox, which demonstrated that becoming part of an identity—the right wing cause—was more profitable than adherence to truth and traditional journalistic standards.

individuals and firms actually have worse information. Just as in Grossman-Stiglitz (1976, 1980), where a more informationally efficient financial market reduces incentives to gather information, so too here, with the net result that there may be a deterioration in the information ecosystem. This may be referred to as the competitive effect of the new technology.

Digital platforms like Google not only divert attention from the legacy media, but steal information from them, information which enhances traffic to them and is accordingly an important contributor to their market value. This may be referred to as the theft effect.[5] The resulting decrease in returns to the production of news leads to a diminution in the supply of information .

1.2 AI

AI puts these perverse incentives on steroids. While AI has the potential to lower information processing and dissemination costs, it undermines the incentives for private producers to acquire and process accurate, timely, and reliable content, as AI firms take information from those who have traditionally been engaged in these activities, such as the legacy media, and combine it with that produced by others. At least, in the case of search engines and social media, there was typically attribution of the source of information, which the platforms claimed (with little evidence) drove more traffic to the legacy media, and thus they got some benefit, even if it did not entail a "fair sharing" of the benefits of the information that the legacy media had produced (Mateen, Tabakovic, Holder, Schiffrin (2023)). But with AI, there may be no clarity about the source of the information, and the synthesis of multiple sources risks reducing the demand for seeing any particular source.[6] Thus, the likelihood that traffic will be driven to the legacy media will be greatly reduced. If AI firms are allowed to appropriate without compensation the information produced, say, by the legacy media, they will obviously not have adequate incentives to produce such information. In that case, in spite of the improvement in technology, again there could be a deterioration in the quality of the information ecosystem. To echo an old expression, "garbage in, garbage out".

AI may also worsen the problem of the quality of information—again leading to a deterioration in the information ecosystem. There is a risk of flooding the information ecosystem with lower-quality, synthetic, or misleading content, effectively polluting it. The potential

[5] We distinguish between these effects partly because there are different policy implications.

[6] Matters may be even worse: AI sometimes attributes to a legacy media things that were not in the media, and which undermine the credibility of the legacy media. The absence of accountability means that the incentives of AI to curb such activities which undermine trust in the legacy media is limited.

adverse consequences are already evident from social media; AI may exacerbate these—again like adding steroids to a problem that we have not yet learned how to manage. There may be a "war" between the ability to detect misleading content and the ability to create such content in ways that deflect AI's ability to detect. If the latter effect dominates, there will be a deterioration of the overall quality of the information ecosystem. And because removing pollution (e.g. through fact checking) is itself a public good, there will be (without public intervention) insufficient expenditures on "clean up," reinforcing the conclusion that, without public intervention, there is a risk of a worsened information ecosystem.

1.3 Outline

This paper is organized as follows. Section 2 develops a model of the information ecosystem. Section 3 examines the effects of AI and digital platforms in an environment where all information is truthful. Section 4 extends the analysis to an environment in which there is also mis/disinformation. Section 5 concludes. We provide an Appendix with all proofs and derivations.

2 A model of the information ecosystem

This section develops a tractable model of producers and consumers of information. We use it to study how digital platforms and the rise of AI is reshaping the information ecosystem.

We broadly define producers as any individual or organization that meets two conditions. The first is that they create truthful or untruthful content that contributes to the stock of information or misinformation, thereby becoming either a public good or a public bad. There are economic incentives to generate untruthful content: producing truthful information is expensive and, if some consumers cannot distinguish between truth and lies, serving this segment with low-cost untruthful information can be profit-maximizing. There are also ideological and strategic incentives: untruthful content can promote certain political or social agendas or discredit competing ones. These non-monetary benefits can be incorporated into a profit function via their monetary equivalent. The second condition is that producers' business models rely on consumer visits to their physical or digital locations. These visits generate revenue through advertisements, subscriptions, or data collection. Traditional producers include newspapers, broadcasting media, and specialized journals, as well as internet websites, educational platforms, and others.

Let $Q^T(t)$ and $Q^L(t)$ be the total amount (stocks) of truthful and untruthful information in the ecosystem at period t. These can be interpreted as the truthful and untruthful information that is in society's "ether". There is a continuum of $n \in [0,1]$ ex-ante identical producers that at time t produce new truthful and untruthful information, $I^T(n,t)$ and $I^L(n,t)$. With the ever evolving economy/society, the value of information accumulated in the past becomes less relevant. For simplicity, we assume that both types of information have the same depreciation rate δ. We can therefore write the dynamics of the stocks of information as follows:

$$\frac{dQ^T(t)}{dt} = I^T(t) - \delta Q^T(t) \tag{1}$$

$$\frac{dQ^L(t)}{dt} = I^L(t) - \delta Q^L(t) \tag{2}$$

where $I^T(t) \equiv \int_0^1 I^T(n,t)dn$ and $I^L(t) \equiv \int_0^1 I^L(n,t)dn$. Implicitly we are assuming that producers create distinct information. Hence there is no redundant information.[7]

We broadly define consumers as individuals or organizations that allocate attention to searching for and using information. Consider a unit mass of consumers. A fraction ω of them use a perfect screening technology at zero cost and can distinguish between the two types of information. We refer to these consumers as informed. The rest of the consumers, which we label uninformed, have no screening technology and cannot distinguish between both types of information. At each time t, each informed and uninformed consumer decides how much time ("minutes") to spend searching for new information. Let $M^I(t)$ and $M^U(t)$ be the minute functions of the representative informed and uninformed consumer, respectively. Then, we assume that:

$$M^I(t) = v^I I^T(t)^\sigma \tag{3}$$

$$M^U(t) = v^U (I^T(t) + I^L(t))^\sigma \tag{4}$$

where $0 < \sigma < 1$ and $v^I = v^I(Q^T(t), Q^L(t), \gamma)$ and $v^U = v^U(Q^T(t), Q^L(t), \gamma)$ are smooth functions that measure the value of new information.[8] At this point, we only assume that these values might be a function of the stocks of truthful and untruthful information and the efficiency of the information ecosystem measured by the parameter γ (which can be thought of as a vector, each element of which describes one aspect of informational efficiency, e.g.

[7] More generally, we can model $I^j(t), j \in \{T, L\}$ as functions of a $I^j(n,t)$ and a measure of differential news, with limiting cases given by: $I^T(t) = \int_0^1 I^T(n,t)dn$ (perfect differentiation) and $I^T(t) = I(n,t)$ (no differentiation).

[8] One can think that each individual has some subset of the available information (or even a garbling of that information) that depends on $Q^T(t), Q^L(t), \gamma$. This is the information upon which an individual can easily draw from and that will determine their minutes function.

processing, transmitting...). The functions $v^I(\cdot)$ and $v^U(\cdot)$ will play a crucial role in the analysis, and we shall say much more about them later. The minutes function of the informed consumers is increasing and concave in $I^T(t)$, but it does not depend on untruthful information $I^L(t)$. Having a perfect screening technology at their disposal allows informed consumers to disregard lies. The minutes function of the uninformed consumers, however, is increasing and concave in $I^T(t) + I^L(t)$. This is a natural assumption since these consumers cannot distinguish between truths and lies.

How are these minutes distributed across producers? Let $a^I(n,t)$ and $a^U(n,t)$ be the share of minutes visiting producers that the representative informed and uninformed consumers spend on producer n. Then, we assume that:

$$a^I(n,t) = \frac{I^T(n,t)}{I^T(t)} \tag{5}$$

$$a^U(n,t) = \frac{I^T(n,t) + I^L(n,t)}{I^T(t) + I^L(t)} \tag{6}$$

In the case of the informed consumers, the share of minutes depends on the size of producer n's flow of truthful information relative to that of other producers. In the case of uninformed consumers, this share depends on the flow of total information relative to that of other producers.[9]

Consumers do not pay for their access to new information. But advertisers pay information producers to post advertisements, knowing that consumers will be exposed to them when searching/processing new information, and this will affect their purchases. Let $1-\lambda$ be the share of the time used to search/process new information in which consumers are directly visiting producers. Assume that the income generated by visits is directly proportional to the time spent in visits. Therefore, the revenue of producer n is given by the visits by informed consumers, $(1-\lambda)\,\omega a^I(n,t)\,M^I(t)$; plus the visits by uninformed consumers, $(1-\lambda)\,(1-\omega)\,a^U(n,t)\,M^U(t)$. This means that we can write the revenue function of producer n as follows:

$$R(n,t) = (1-\lambda)\left[\omega v^I I^T(t)^{\sigma-1} I^T(n,t) + (1-\omega) v^U \left(I^T(t)+I^L(t)\right)^{\sigma-1}\left(I^T(n,t)+I^L(n,t)\right)\right] \tag{7}$$

[9] More generally, we could allow non-proportional targeting of consumers (opinion leaders, high-centrality nodes), so producers choose whom to reach, not just how much to produce, mapping directly to strategic competition among influencers (Li and Tan (2025)).

The marginal effects on revenue of both types of information are positive and depend on the choices of other producers: $I^T(n',t)$ and $I^L(n',t)$ for $n' \neq n$. Since $\sigma < 1$, the marginal revenue of both types of information is decreasing on the information produced by other producers. Thus, on the marginal revenue side, producer choices are strategic substitutes: as the flow of information of other producers grows, the incentives to produce information declines.[10]

Lies are cheaper to produce than the truth. In particular, we assume that:

$$C\left(n,t\right) = \frac{c}{\varepsilon}\left[I^T\left(n,t\right) + \theta I^L\left(n,t\right)\right]^{\varepsilon} \tag{8}$$

where $\varepsilon > 1$, $0 < \theta < 1$ and $c = c(Q^T(t), Q^L(t), \gamma)$ is a smooth function that measures the cost of producing new information. At this point, we only assume that it might depend on the stocks of truthful and untruthful information and the efficiency of the information ecosystem. We shall say much more about this function later. The parameter θ, measures the cost of producing untruthful information relative to truthful information. Since $0 < \theta < 1$, lies are always cheaper to produce than the truth. As $\theta \to 1$,lies become arbitrarily as expensive as the truth. As $\theta \to 0$, lies become arbitrarily cheap. The cost of producing new information is increasing and convex in both flows of information, and it is not affected by the choices of other producers. The latter means that, on the cost side, producer actions are neither substitutes nor complements from a strategic viewpoint.[11]

Summing up, producers maximize the net present value of profits by choosing how much truthful and untruthful information to produce, $I^T(n,t)$ and $I^L(n,t)$. We have made two assumptions that simplify the analysis dramatically. First, future profits of producers do not depend on the flow of information that they create today, $I^T\left(n,t\right)$ and $I^L\left(n,t\right)$.[12] Second, producers are atomistic and they take as given the aggregate stocks of information, $Q^T(t)$ and $Q^L\left(t\right)$. Jointly, these assumptions imply that individual producer choices at time t have no effects on their future profits and, as a result, their dynamic maximization problem breaks

[10]It would be interesting to explore network effects and other forces that can make producer actions strategic complements. We leave this for future research.

[11]In a richer model, an increase in the flow of information by other producers raises the prices of inputs (including labor) and, as a result, the costs of producer n. Thus, we would have that producer actions are also strategic substitutes on the cost side.

[12]This would not be the case if producers had reputation costs. In making their decisions about the flow of new information, they would take into account effects at later dates. This might turn out to be an important extension. An interesting recent paper provides field evidence showing that greater concern over AI-generated misinformation reduces overall trust in news content but increases demand of most trusted sources (Campante, Durante, Hagemeister, Sen 2025).

down into a sequence of independent static problems. Thus, the flow of information created in period t is the outcome of a simple static producer game.

What types of Nash equilibria arise in this producer game? Let $P(t)$, be the share of untruthful information created in period t, that is, $P(t) \equiv \frac{I^L(t)}{I^T(t) + I^L(t)}$. This share can be interpreted as a measure of how polarizing new information is.[13] It turns out that the information ecosystem induces two regimes. For some parameter values, only truthful information is produced ($P(t) = 0$). In this case, we say we have a truthful equilibrium. For the remaining parameter values, both truthful and untruthful information are produced ($P(t) > 0$). We call this the truth-lies equilibrium. Within this second regime, there are some limiting cases in which the share of lies approaches one. This can occur either because no truthful information is produced or because the production of lies grows without bound. We now examine both regimes separately.

3 Truthful equilibrium

We start by analyzing the truthful equilibrium, which serves as a benchmark on how AI and digital platforms would affect the information ecosystem in a world without misinformation.[14]

Proposition 1 (Truthful equilibrium) Suppose $\frac{\omega v^I}{(1-\omega) v^U} \frac{\theta}{1-\theta} > 1$. Then the *unique* Nash equilibrium is such that $P(t) = 0$, and:

$$I^T(t) = \left((1-\lambda) \frac{\omega v^I + (1-\omega) v^U}{c} \right)^{\frac{1}{\varepsilon - \sigma}}$$

$$I^L(t) = 0$$

We attain a truthful equilibrium when: informed consumer demand for information is high (high ωv^I), uninformed consumer demand for information is low (low $(1-\omega) v^U$) and/or lies are not that cheap relative to truths (high θ). In such an environment, all producers will want to solely produce truthful information. The flow of new information $I^T(t)$ is high when: many consumers obtain their information directly from producers (low λ), the value of information is high (high $\omega v^I + (1-\omega) v^U$) and the cost of information is low (low c).

[13] If $P(t) = 0$, the flow is not polarizing because all consumers receive the same information. If $P(t) = 1$, the flow's polarizing effect is maximized as informed and uninformed consumers receive totally different information at time t.

[14] The Appendix contains the proofs of all propositions.

If the information ecosystem starts with a positive stock of lies, these lies depreciate over time and we reach a steady state with $Q^L(t) = 0$. Thus, it makes sense to simplify the analysis by assuming that information ecosystem already starts with $Q^L(t) = 0$. Hence, in the remainder of this section, whenever we say "information" without making any distinction, we are referring to "truthful information". It also makes sense that, in an environment in which there is only truthful information, the value of information is the same for informed and uninformed consumers. Thus, in the remainder of this section, we assume that $v^I = v^U = v$.

3.1 Dynamics and equilibrium

With these simplifications, we can now write the dynamics of the stock of information as follows:

$$\frac{dQ^T(t)}{dt} = \left((1-\lambda)\frac{v\left(Q^T(t),\gamma\right)}{c\left(Q^T(t),\gamma\right)}\right)^{\frac{1}{\varepsilon-\sigma}} - \delta Q^T(t) \tag{9}$$

where now we write the functions describing the value and cost of new information explicitly as functions of the stock of information, $Q^T(t)$, and the efficiency of the system, γ. Up to now, we have not made assumptions about these functions because these were not needed. But now we are forced to make assumptions about them, because the dynamics of the system crucially depend on the ratio of value to cost of information.

The first assumption we make is that $\lim_{Q^T(t)\to 0} \frac{v\left(Q^T(t),\gamma\right)}{c\left(Q^T(t),\gamma\right)} = 0$. The cost of information grows without bound as the stock of information goes to zero. One cannot simply produce information without any knowledge. We assume that even if the value of information increases, it does so less rapidly. This assumption implies that the information ecosystem always has a steady state with zero information (information collapse).

The second assumption is that $\lim_{Q^T(t)\to\infty} \frac{v\left(Q^T(t),\gamma\right)}{c\left(Q^T(t),\gamma\right)} = 0$. The value of new information shrinks to zero as the stock of information goes to infinity. One does not need more additional information when one's knowledge is unbounded. The same does not happen with the cost of information. This assumption ensures that the stock of information does not grow without a bound.

The number of steady states and their location depends on the behavior of the ratio of value to cost of information. Our assumptions ensure that this ratio starts at zero when

$Q^T(t) = 0$ and eventually returns to zero when $Q^T(t) = \infty$. What happens in the middle depends on the relative growth in the value and cost of information as $Q^T(t)$ grows. We assume that initially $\frac{v_Q}{v} > \frac{c_Q}{c}$ and the ratio increases. The value of information rises fast relative to its cost because, if the stock of information is low, additions to this stock raise the marginal value of new information (Radner and Stiglitz (1984)) and lowers marginal costs by improving the information acquisition technology. At some point the growth in this ratio starts to decline and eventually reverses $\frac{v_Q}{v} < \frac{c_Q}{c}$. The cost of information now rises fast relative to its value because there are diminishing returns to information and there are less "low-hanging fruits" available for producers to find.

To sum up, our assumptions ensure that the ratio of value to costs exhibits an inverse-U shape starting at zero when the stock of information goes to zero, and returning to zero as the stock of information goes to infinity. Our final assumption is that there exists a value Q such that:

$$\left((1-\lambda)\frac{v(Q,\gamma)}{c(Q,\gamma)}\right)^{\frac{1}{\varepsilon-\sigma}} > \delta Q \tag{10}$$

This additional assumption ensures that there is always a steady state besides the information collapse. Moreover, this steady state is stable. We refer to Condition (10) as the " no information collapse" condition. If the ratio of value to cost of information grows slowly when $Q^T(t) = 0$ (at a rate lower than $\varepsilon - \sigma$), there are an odd number of steady states, and the information collapse is a stable equilibrium. This is the case we focus from now on (to illustrate we focus on the case of three steady states).

Figure 1 shows the dynamics of the information ecosystem.[15] If the initial stock of information is high, the information ecosystem converges to the highest steady state (from now on referred to as the equilibrium and denoted as Q^{T*}). If the initial stock of information is low, the ecosystem experiences an information collapse. The model shows that there is path dependence, that is, even temporary aberrations in the information ecosystem (a sufficiently large drop/increase in $Q^T(t)$) can have permanent consequences.

[15] In drawing the figure, we calibrate equation (9) as follows: $\varepsilon = 1.5$, $\sigma = 0.5$, $\delta = 0.75$, $\lambda = 0.2$, $\gamma = 1$, $\frac{v(Q^T(t),\gamma)}{c(Q^T(t),\gamma)} = \frac{v(\gamma Q^T(t))}{c(\gamma Q^T(t))} = \max\{0, \gamma Q^T(t)^3 e^{-0.6\gamma Q^T(t)}\}$.

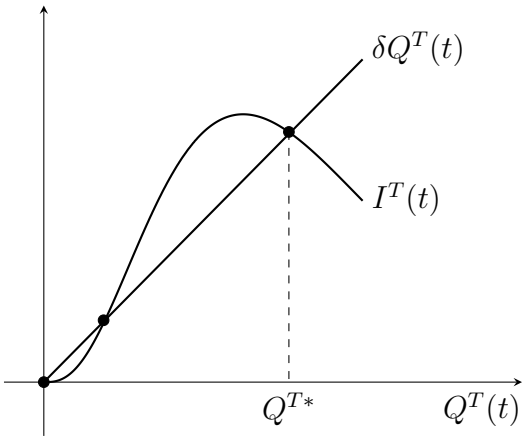


Figure 1: Equilibrium level of Q^{T*}.

3.2 Welfare and comparative statics

There is no presumption that the steady state stock of information is socially optimal. But it would be inaccurate to measure welfare simply as the total stock of information in society, $Q^T(t)$, given that we would not be taking into account the efficiency of the information ecosystem. More precise is to measure welfare as some function $q(t) = q\left(Q^T(t), \gamma\right)$ which could be interpreted as the personal information set of the representative consumer. That is, the information upon which an individual can easily draw from.

Our enquiry here, is: how will AI affect the information ecosystem. AI improves the ways in which we acquire, process, and transmit information produced by others. This can be modeled as how an increase in γ affects $q\left(Q^T(t), \gamma\right)$. At the same time, if consumers spend more of their time searching/processing new information through AI intermediaries, producers of information will obtain less direct visits, endangering their business model. This can be modeled as an increase in λ.[16] We proceed to study each of these effects separately.[17]

[16]For purposes of most of this section, digital platforms (traditional information aggregators) have similar qualitative consequences as that of AI, although quantitatively these will differ. AI is likely to increase γ by a larger magnitude than digital platforms as it operates across a far broader set of disciplines and draws on a more diverse range of informational inputs in the process of generating and disseminating information. AI is also likely to increase λ by a greater magnitude than digital platforms since they offer specific and detailed information and obscure attribution more thoroughly. 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 (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)).

[17]We are simply considering the impact of AI as a technological shock. Although we will not go down this path, further insights can be gained by adding AI firms as players who choose γ and λ to maximize

To determine the effects of an improvement in efficiency (increase in γ), we need to make assumptions about the effects of γ on the ratio of the value to cost of information. Here, we have two opposing forces. On one hand, an increase in γ reduces the value of new information. If information is transmitted faster (more efficiently) throughout society, the value of obtaining new information before others is lower (Alchian and Demsetz (1972)). This is analogous to how (it is alleged that) weaker intellectual property rights weaken incentives to innovate or how information efficiency of financial markets lowers incentives to acquire information. On the other hand, an increase in γ reduces the cost of new information. It is reasonable to assume that higher efficiency in managing information lowers production costs. Thus, the effect of γ on the ratio of value to cost could go either way. Naturally, an increase in γ could lead to greater information production and an overall improvement of the information ecosystem. But it could also lead to a worsening of the information ecosystem if information production falls. It is hard to tell which force will predominate in this setting.

Figure 2 illustrates this point with an example in which the force that predominates depends on $Q^T(t)$. In particular, we assume that $v = v\left(\gamma Q^T(t)\right)$ and $c = c\left(\gamma Q^T(t)\right)$.[18] Hence, an increase in γ shifts the curve $I^T(t)$ upwards for small $Q^T(t)$ but downwards for large $Q^T(t)$. Interestingly, this case leads to a reduction in the steady state stock of information. This dynamic mirrors the logic of the Grossman-Stiglitz paradox (1980).[19]

discounted profits. In such an extension, the nature of competition among AI firms would be an important determinant of the effects of AI on the information ecosystem. For instance, a monopolist would strategically decide how much consumer attention/information to steal from producers, taking into account that too much stealing eliminates the incentives to produce information in the future. But competitive AI firms might not take this into account and lead to an information collapse. We leave for further research the study of the optimal market structure of AI firms.

[18] We set $\gamma_L = 1$ and $\gamma_H = 1.4$. All other parameters are calibrated as in Figure 1.

[19] If markets (or AI-intermediaries) were perfectly informative (fully informationally efficient markets, in the traditional sense (see, e.g. Fama (1970)), then no agent would have an incentive to acquire or produce costly information. Markets would be fully efficient in transmitting the information that exists within the information ecosystem, but the only information within the system would be that which could be obtained costlessly -with appropriate normalization, we can think of as a zero information system - and therefore a zero knowledge ecosystem.

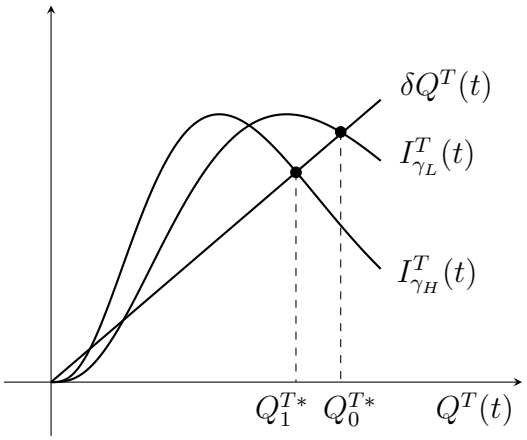


Figure 2: An increase in γ can lower Q^{T*}.

Does this mean that welfare declines as a result of AI? Let us assume that $q(t) = \gamma Q^T(t)$. In this case, the effect of AI on welfare is ambiguous. On one hand, γ increases. On the other hand, $Q^T(t)$ declines. A simple way to think about this result is that the total stock of information in society, $Q^T(t)$, is a "pie" and γ is the share of the pie that any individual in society has access to. An increase in γ has the immediate effect of increasing the share of the pie. But it then slowly reduces the size of the pie as the economy transitions to the new steady state. Whether ex-post knowledge is greater or lower than ex-ante is therefore unclear. Figure 3 illustrates the case where the fall in the size of the pie predominates and welfare drops (the opposite is, of course, possible too). Point A is the ex-ante equilibrium, associated with a low level of γ. When γ increases, we immediately jump to point B (the share of the " stock of information"pie" that gets translated into " meaningful knowledge" increases). Everyone praises the transformative role of the new technology. However, as $Q^T(t)$ falls and we move towards point C, societal welfare declines.

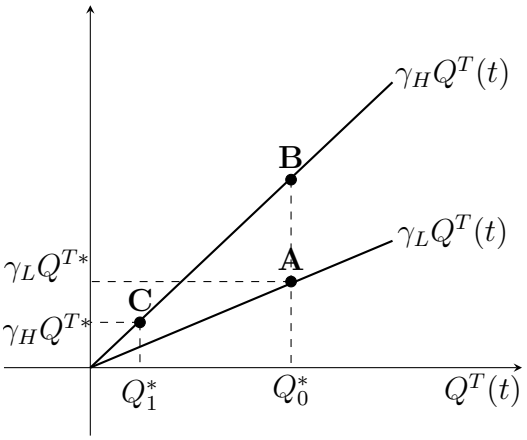


Figure 3: An increase in γ can decrease γQ^{T*}.

It is much easier to determine the effects of a deterioration of the business model (increase in λ). An increase in λ unambiguously reduces the incentives to produce new information and the stock of information declines. Figure 4 illustrates this.[20] It also follows that societal welfare unambiguously declines.

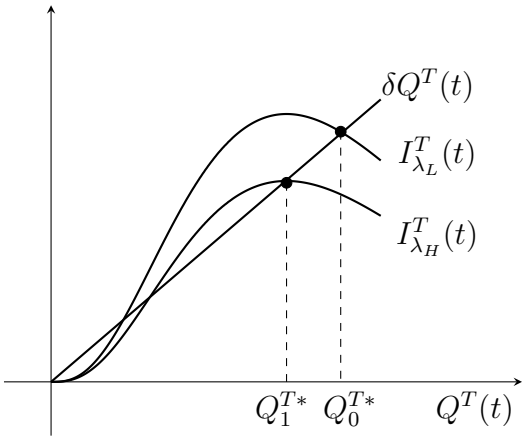


Figure 4: An increase in λ lowers Q^{T*}.

3.3 Information collapse

Finally, we note that, by increasing γ and λ, AI might lead to the violation of Condition (10), posing an existential threat to the information ecosystem. At the current state of AI, hallucinations and inaccuracies serve as a natural brake on user substitution away from primary sources. Because users cannot fully rely on AI outputs, they continue to engage with information producers preserving the incentives for information production. However, as AI intermediaries become increasingly accurate and contextually fluent, with the ability to automatically check and verify the original sources to which human actors refer, the need to consult original sources diminishes. In the limit, if AI becomes perfectly reliable-or reliable enough and all information consumption is mediated by such systems, incentives to produce new information collapse. No one invests in producing accurate information when their work is instantly absorbed and intermediated by an AI that captures all downstream attention and value.[21]

[20] We set $\lambda_L = 0.2$ and $\lambda_H = 0.4$. All other parameters are calibrated as in Figure 1.

[21] Much the same applies to the deeper problem of information aggregation and processing. While the platforms champion their role in information transmission, AI does that as well as acquiring and transmitting information produced by others. Grossman and Stiglitz (1976) showed that (futures) markets, on their own, do not do a good job in information aggregation, contrary to the assertions of Hayek (1944) and others. Forecasting accuracy can be increased when such information is combined with other information, but if that kind of analysis is costly, and if the fruits of that kind of analysis can be obtained at a low enough cost, then there will be diminished incentives to acquire the additional information and/or use such additional

4 Truth-lies equilibrium

In this section we study the presence of mis/disinformation. We start with the following proposition:

Proposition 2 (Truth-lies equilibrium) Suppose $\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)} \leq 1$. Then the *unique* producer Nash Equilibrium is such that $P(t) = 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}} \geq 0$, and:

$$I^T(t) = (1-P(t))\left[(1-\lambda)\frac{\omega v^I(1-P(t))^{\sigma-1}+(1-\omega)v^U}{c(1-P(t)(1-\theta))^{\varepsilon-1}}\right]^{\frac{1}{\varepsilon-\sigma}}$$

$$I^L(t) = P(t)\left[(1-\lambda)\frac{\omega v^I(1-P(t))^{\sigma-1}+(1-\omega)v^U}{c(1-P(t)(1-\theta))^{\varepsilon-1}}\right]^{\frac{1}{\varepsilon-\sigma}}$$

We can uniquely pin down $P(t)$. If the degree of polarization in the information ecosystem, $P(t)$, is high enough, all producers prefer to only produce truthful information. The reason is that truthful information is scarce and its value (in terms of minutes) is high. Hence $P(t)$ falls as a result. The opposite occurs if $P(t)$ is low enough. In this case, untruthful information is relatively scarce and all producers prefer to only produce untruthful information thereby increasing $P(t)$. These stabilizing dynamics leads the information ecosystem towards a unique $P(t)$. This also pins down the flows of truthful and untruthful information.[22]

The flows of new information $I^T(t)$ and $I^L(t)$ are high when: many consumers obtain their information directly from producers (low λ), the value of information is high (high $\omega v^I(1-P(t))^{\sigma-1}+(1-\omega)v^U$) and the cost of information is low (low $c[1-P(t)(1-\theta)]^{\varepsilon-1}$). Relative to the previous regime, we note three differences. The first is that a fraction of the new informational flows consist of lies. The second is that the value of information as perceived by consumers is lower, as informed consumers disregard lies. The third is that the cost of information is lower because producing lies is cheaper.

Now, the dynamics of the information ecosystem can be obtained by substituting the expressions in Proposition 2 into the laws of motion in Equations (1) and (2). This is a complicated

information in combination with the future price information to make better forecasts. Again, the amount of usable information within society is diminished.

[22] We cannot pin down individual producer flows $I^T(n,t)$ and $I^L(n,t)$, because in equilibrium producers are indifferent between producing truths or lies. Thus, the model is not inconsistent with producers specializing in different types of information ex-post even though they are identical ex-ante. For more details, see the Appendix.

dynamical system, and a full analysis of it is not possible in such a short paper. But we can derive most of the relevant insights using a simple case. Assume that $v^I = v^U = v$, and let v and c be only a function of $Q^T(t)$ and γ, but not a function of $Q^L(t)$. Admittedly, our main justification for these assumptions is that they provide a tractable benchmark. But one could also justify them by arguing that lies are short-lived in the sense that, even though uninformed consumers cannot recognize lies on impact, they are able to do so without too much delay. In this case, uninformed consumers cannot distinguish between flows of truthful and untruthful information, but they can distinguish between stocks of truthful and untruthful information.

Under these assumptions, we have that:

$$P(t) = P \equiv 1 - \left(\frac{\omega}{1-\omega}\frac{\theta}{1-\theta}\right)^{\frac{1}{1-\sigma}} \tag{11}$$

$$\frac{dQ^T(t)}{dt} = (1-P)\left((1-\lambda)\frac{v\left(Q^T(t),\gamma\right)\left[\omega(1-P)^{\sigma-1}+1-\omega\right]}{c\left(Q^T(t),\gamma\right)\left[1-P(1-\theta)\right]^{\varepsilon-1}}\right)^{\frac{1}{\varepsilon-\sigma}} - \delta Q^T(t) \tag{12}$$

$$\frac{dQ^L(t)}{dt} = P\left((1-\lambda)\frac{v\left(Q^T(t),\gamma\right)\left[\omega(1-P)^{\sigma-1}+1-\omega\right]}{c\left(Q^T(t),\gamma\right)\left[1-P(1-\theta)\right]^{\varepsilon-1}}\right)^{\frac{1}{\varepsilon-\sigma}} - \delta Q^L(t) \tag{13}$$

This benchmark case is tractable for two reasons. The first is that the share of lies in the new flows of information is constant over time $P(t) = P$. Thus, in any non-zero steady state, the share of lies in the information ecosystem converges to P, that is, $\frac{Q^L(t)}{Q^T(t)+Q^L(t)} \to P$. This means that, if we know the evolution of the stock of truthful information, we also know the evolution of the stock of lies. The second reason is because the evolution of the stock of truthful information is independent of the stock of lies in the system, that is, the stock of lies does not appear in Equation (12). This means that we can study the evolution of the stock of truthful information using a simple one-dimensional dynamical system.

The properties of this system depend again on the assumptions we make about the ratio of the value to the cost of information. We adopt the assumptions of the previous section, except that we now re-write Condition (10) by saying that there exists a value Q such that:

$$(1-P)\left((1-\lambda)\frac{v(Q,\gamma)\left[\omega(1-P)^{\sigma-1}+1-\omega\right]}{c(Q,\gamma)\left[1-P(1-\theta)\right]^{\varepsilon-1}}\right)^{\frac{1}{\varepsilon-\sigma}} > \delta Q \tag{14}$$

If this condition holds, there are non-zero steady states for the stock of truthful information, and the dynamics of $Q^T(t)$ are qualitatively similar to those shown in Figure 1. If this condition fails, there is only one steady state which has zero truthful information. We saw in the previous section that an increase in γ and/or λ could lead the ecosystem to zero information steady state. An interesting novelty here is that a reduction in ω and/or θ can also generate a steady state without truthful information as P increases.

So what is new here about the effects of AI and digital platforms? As discussed in the previous section, these technologies improve efficiency (raise γ) and deteriorate the business model of producers (raise λ). All the effects discussed in the previous section apply also here. The presence of mis/disinformation in the ecosystem does not affect this part of the analysis. But the presence of mis/disinformation creates two additional channels through which AI and digital platforms might affect the information ecosystem: the ability of consumers to detect lies (ω) and the relative cost of producing lies (θ). These are key determinants of polarization and the amount of mis/disinformation in the system. We focus on these new channels next.

AI has an ambiguous effect on the proportion of informed/uninformed consumers (ω). On the one hand, AI can be used by producers to make untruthful information harder to detect. On the other hand, AI can also be used by consumers to improve detection of untruthful information. We refer to these conflicting effects of AI on ω as the "drone war effect". Which force predominates? Although this is an open question, there are reasons to believe that the the first force dominates. It is undeniable that AI is helping producers create untruthful information (creation of malicious bots, propagation of "fake news", targeted and personalized news feed algorithms...). However, the extent to which AI is currently helping consumers detect untruthful information is unclear. AI-intermediaries can hallucinate, and more importantly, when users search for information they are not constantly verifying everything they consume with AI. So, we presume that the effect of AI on ω is likely negative. This leads to an increase in polarization and a reduction in the steady state stock of truthful information.

The effect of AI on relative cost of producing lies (θ), however, is more convincingly unambiguous: AI reduces the cost of producing untruthful information relative to the cost of producing truthful information, effectively making it very cheap to produce lies. This also leads to an increase in polarization and a reduction in the steady state stock of truthful information. As we have just discussed, this could even lead to a steady state with no truthful information. There is mounting evidence (AI-generated "news" sites, social media clickbait

images, "AI slop"...) that unless producers are held accountable for producing untruthful information, we will be heading in this direction.

Digital platforms also incentivize the production of mis/disinformation and increase polarization. But the channels are different: digital platforms do not affect ω and θ directly. Despite this, they might generate effects that could be interpreted along similar lines. Although not modeled here, both AI and digital platforms reduce producer accountability by obscuring attribution. Hence, digital platforms do indirectly reduce the relative cost of lies (but certainly not as much as AI). Suppose also that consumers are especially engaged by extreme events or information that enrages them ("engagement through enragement"). This can incentivize producers to create this type of information, even if untruthful, to satisfy the demand from digital platforms. Also, digital platforms such as social media target consumers individually and give different information to different consumer types, increasing polarization. This is not the case of AI-intermediaries who, for the moment, generate synthesis and a more neutral narrative giving each consumer similar information, reducing polarization.

5 Conclusions

This paper is a first attempt to formalize the analysis of the consequences of digital platforms and AI on the information ecosystem. We have questioned the nirvana promised by the champions of these new technologies, if left to their own devices: competition in this arena does not necessarily lead to the well-being of society. We have provided a variety of reasons that the quality of the information ecosystem may deteriorate—even as these innovations increase the ability to process information and decrease the costs of disseminating information. In addition, our framework opens up a new research agenda on measurement, aimed at quantifying the key channels of our model, such as the relative cost of producing lies (θ), the fraction of informed consumers (ω), and the share of advertising revenue captured by intermediaries (λ). Finally, we emphasize that an "information collapse" does not require the advent of transformative AI (TAI); it is enough for AI to be merely "good enough" at synthesizing existing information. This means that the period between today's systems and the potential arrival of TAI is itself fraught with risk, as AI platforms can undermine incentives for information production without yet replacing it.

Policy responses

The design of appropriate intellectual property, accountability, competition and regulatory

frameworks will be critical in mitigating the adverse effects on the information ecosystem we've described. Our legal frameworks clearly need to adapt to the challenges posed by these new technologies. For instance, in the absence of accountability, there is little incentive for digital platforms not to circulate mis and disinformation, and if such mis and disinformation garners more eyeballs, there are even incentives for expanding the reach of such information. The willingness of AI firms to share their profits with those who produce the information they rely on may be affected by the competitive structure in AI itself, and by the terms of intellectual property legislation. Current frameworks were obviously created before these challenges presented by AI and the digital platforms were on the scene. Critical constructs like fair use were not designed to address whether such use of others' intellectual property would increase or decrease the quality of the information ecosystem in the presence of AI or the digital platforms, and courts that have naively extended the reach of such constructs may be imposing real harms on our society. A full analysis would, however, take us beyond the scope of this paper.

References

[1] Alchian, A. A., and H. Demsetz. 1972. "Production, Information Costs, and Economic Organization." *American Economic Review* 62(5): 777–95.

[2] Arrow, Kenneth J. 1962. "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, NJ: Princeton University Press.

[3] Calzada, J., and R. Gil. 2020. "What Do News Aggregators Do? Evidence from Google News in Spain and Germany." *Marketing Science* 39(1): 134–67.

[4] Campante, Filipe R., Ruben Durante, Felix Hagemeister, and Ananya Sen. 2025. "GenAI Misinformation, Trust, and News Consumption: Evidence from a Field Experiment." NBER Working Paper No. 34100.

[5] Fama, Eugene F. 1970. "Efficient Capital Markets: A Review of Theory and Empirical Work." *Journal of Finance* 25(2): 383–417.

[6] Fama, Eugene F. 1991. "Efficient Capital Markets: II." *Journal of Finance* 46(5): 1575–1617.

[7] Greenwald, B., and J. E. Stiglitz. 1992. "Information, Finance and Markets: The Architecture of Allocative Mechanisms." *Industrial and Corporate Change* 1(1): 37–63.

[8] Grossman, S., and J. E. Stiglitz. 1976. "Information and Competitive Price Systems." *American Economic Review* 66(2): 246–53.

[9] Grossman, S., and J. E. Stiglitz. 1980. "On the Impossibility of Informationally Efficient Markets." *American Economic Review* 70(3): 393–408.

[10] Jeon, D.-S., and N. Nasr. 2016. "News Aggregators and Competition among Newspapers on the Internet." *American Economic Journal: Microeconomics* 8(4): 91–114.

[11] Li, Wei, and Xu Tan. 2025. "Dynamic Network Influence: The Art of Strategic Messaging." SSRN Working Paper No. 5066467.

[12] Lyu, Liang; James Siderius; Hannah Li; Daron Acemoglu; Daniel Huttenlocher; and Asuman Ozdaglar. 2025. "Wikipedia Contributions in the Wake of ChatGPT." WWW '25.

[13] Mateen, H., Tabakovic, H., Holder, P., Schiffrin, A. 2023. Paying for News: What Google and Meta Owe U.S. Publishers. SSRN / IPD Columbia (Working paper).

[14] Stiglitz, J. E. 1975a. "Information and Economic Analysis." In J. M. Parkin and A. R. Nobay, eds., *Current Economic Problems*. Cambridge: Cambridge University Press, 27–52.

[15] Stiglitz, J. E. 1975b. "The Theory of Screening, Education and the Distribution of Income." *American Economic Review* 65(3): 283–300.

[16] Stiglitz, J. E. 1987. "On the Microeconomics of Technical Progress." In Jorge M. Katz, ed., *Technology Generation in Latin American Manufacturing Industries*. New York: St. Martin's Press, 56–77.

[17] Stiglitz, J. E. 1989. "Using Tax Policy to Curb Speculative Short-Term Trading." *Journal of Financial Services Research* 3(2–3): 101–15.

[18] Stiglitz, J. E. 1999. "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*. New York: Oxford University Press, 308–25.

[19] Stiglitz, J. E. 2014. "Intellectual Property Rights, the Pool of Knowledge, and Innovation." NBER Working Paper No. 20014.

[20] Stiglitz, J. E. 2021. "The Media: Information as a Public Good." Paper presented to the Pontifical Academy of Social Sciences.

[21] Athey, S., M. Möbius, and J. Pal. 2021. "The Impact of Aggregators on Internet News Consumption." NBER Working Paper No. 28746.

Appendix

This Appendix proves Proposition 1 and Proposition 2.

Profit maximization by producer n implies that:

$$(1-\lambda)\left[\omega v^I I^T(t)^{\sigma-1} + (1-\omega)\, v^U \left[I^T(t) + I^L(t)\right]^{\sigma-1}\right] \leq c\left[I^T(n,t) + \theta I^L(n,t)\right]^{\varepsilon-1} \quad (15)$$

$$(1-\lambda)(1-\omega)\, v^U \left[I^T(t) + I^L(t)\right]^{\sigma-1} \leq \theta c\left[I^T(n,t) + \theta I^L(n,t)\right]^{\varepsilon-1} \quad (16)$$

Conditions (15) and (16) state that profit-maximizing choices must be such that the marginal revenue of additional units of truthful or untruthful information cannot exceed their marginal cost. If Condition (15) holds as a strict inequality, then $I^T(n,t) = 0$. If Condition (16) holds as a strict inequality, then $I^L(n,t) = 0$.

Proposition 1

Suppose $\dfrac{\omega v^I}{(1-\omega)\, v^U}\dfrac{\theta}{1-\theta} > 1$. Then the *unique* Nash equilibrium is such that $P(t) = 0$, and:

$$I^T(t) = \left((1-\lambda)\frac{\omega v^I + (1-\omega)\, v^U}{c}\right)^{\frac{1}{\varepsilon-\sigma}}$$

$$I^L(t) = 0$$

Proof:

Suppose $\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)} > 1$.

First, we will establish that $I^L(n,t) = 0, \forall n,t$. Suppose, by contradiction, there is a producer n such that $I^L(n,t) > 0$. Then, since producer n is maximizing, it must be that condition (16) holds with equality. Then, substituting condition (16) into condition (15), and rearranging the expression in terms of $P(t)$, we find that:

$$\frac{\omega v^I \theta}{(1-\omega)v^U(1-\theta)} \leq (1-P(t))^{1-\sigma}$$

which is a contradiction since $P(t) \in [0,1]$ and $\sigma < 1$.

This implies that $I^T(n,t) > 0, \forall n,t$. This is because the costs are convex and start at 0 while producers have a positive and constant marginal revenue. Therefore, it will always generate some positive profit to produce some $I^T(n,t)$. Hence, condition (15) holds with equality. Thus, given our assumptions, the unique equilibrium is for all producers to follow the same strategy: $I^T(n,t) = I^T(t)$ and $I^L(n,t) = I^L(t)$ for all $n \in [0,1]$, and the producer game has a unique Nash equilibrium given by:

$$I^T(t) = \left[\frac{(1-\lambda)}{c}(\omega v^I + (1-\omega)v^U)\right]^{\frac{1}{\varepsilon-\sigma}} \tag{17}$$

$$I^L(t) = 0 \tag{18}$$

□

Proposition 2

Suppose $\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)} \leq 1$. Then the *unique* producer Nash Equilibrium is such that $P(t) = 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$, and:

$$I^T(t) = (1-P(t))\left[(1-\lambda)\frac{\omega v^I(1-P(t))^{\sigma-1} + (1-\omega)v^U}{c(1-P(t)(1-\theta))^{\varepsilon-1}}\right]^{\frac{1}{\varepsilon-\sigma}}$$

$$I^L(t) = P(t)\left[(1-\lambda)\frac{\omega v^I(1-P(t))^{\sigma-1} + (1-\omega)v^U}{c(1-P(t)(1-\theta))^{\varepsilon-1}}\right]^{\frac{1}{\varepsilon-\sigma}}$$

Proof:

Suppose $\frac{v^I}{v^U}\frac{\omega\theta}{(1-\omega)(1-\theta)} \leq 1$.

First, we will establish that $P(t) = 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$ by discarding the alternatives:

i) Suppose, by contradiction, that $P(t) > 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$. This implies that $I^L(n,t) = 0, \forall n,t$. Suppose not, and let $I^L(n,t) > 0$. Then, condition (16) holds with equality. Then,

substituting condition (16) into condition (15) we find that:

$$P(t) \leq 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$$

which is a contradiction. But if $I^L(n,t) = 0, \forall n, t$ we then have that $P(t) = 0$ which is a contradiction.

ii) Suppose, by contradiction, that $P(t) < 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$. This implies that $I^T(n,t) = 0, \forall n, t$. Suppose not, and let $I^T(n,t) > 0$. Then, condition (15) holds with equality. Then, substituting condition (15) into condition (16) we find that:

$$P(t) \geq 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$$

which is a contradiction. But if $I^T(n,t) = 0, \forall n, t$ we then have that $P(t) = 1$ which is a contradiction.

Let us now consider the case in which $P(t) = 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}}$. Since $0 < P(t) < 1$, we now have that both (15) and (16) hold with equality. Then, (15) implies that each producer n is producing a combination of $I^T(n,t) \geq 0$ or $I^L(n,t) \geq 0$ such that:

$$I^T(n,t) + \theta I^L(n,t) = \left(\frac{(1-\lambda)}{c}\left[\omega v^I I^T(t)^{\sigma-1} + (1-\omega)\, v^U (I^T(t) + I^L(t))^{\sigma-1}\right]\right)^{\frac{1}{\varepsilon-1}} \quad (19)$$

That is, if the aggregate degree of polarization is exactly as assumed, producer n is indifferent about producing truthful information or lies, provided that the total flows satisfy equation (19).

Since individual productions are not determined, there are many Nash equilibria for the producer game. Each of them correspond to a different distribution of flows among producers. But in all these equilibria, aggregate flows of information must be such that they satisfy equation (19) and:

$$\frac{I^L(t)}{I^T(t) + I^L(t)} = 1 - \left[\frac{\omega v^I}{(1-\omega)v^U}\frac{\theta}{(1-\theta)}\right]^{\frac{1}{1-\sigma}} \quad (20)$$

We now use (19) and (20) to find to find $I^T(t) + I^L(t)$ and express $I^T(t)$ and $I^L(t)$ as fractions of the total sum. Thus, it follows that, if $\frac{v^I}{v^U}\frac{\omega\theta}{(1-\omega)(1-\theta)} \leq 1$, the equilibrium flows of truthful and untruthful information are given by:

$$I^T(t) = (1 - P(t)) \left[(1-\lambda) \frac{\omega v^I (1 - P(t))^{\sigma-1} + (1-\omega) v^U}{c(1 - P(t)(1-\theta))^{\varepsilon-1}} \right]^{\frac{1}{\varepsilon-\sigma}}$$

$$I^L(t) = P(t) \left[(1-\lambda) \frac{\omega v^I (1 - P(t))^{\sigma-1} + (1-\omega) v^U}{c(1 - P(t)(1-\theta))^{\varepsilon-1}} \right]^{\frac{1}{\varepsilon-\sigma}}$$

□