NETWORK LAW REVIEW

ARTICLE

The Matthew Effect at Scale: Attention Scarcity and the AI Output Explosion

Thibault Schrepel

Abstract: Generative AI is lowering the marginal cost of academic writing. A natural prediction is that this will allow more researchers to exercise scientific influence. This article argues that the effects are more complex. An output explosion amplifies the Matthew effect, concentrates reputational gains among established scholars, and contributes to the emergence of increasingly stratified publication tracks that rarely intersect.

*

1. Why This Belongs Here

The Network Law Review takes networks seriously. Academic publishing is a network. Authors, journals, citations, and editorial boards form a system of nodes and edges in which reputation travels along well-worn paths. When a new technology alters the cost of producing the nodes, the structure shifts. Generative AI is that technology. This short article examines what happens to academic publishing networks when the marginal cost of writing a paper reduces.

The argument runs in four steps. Generative AI will produce a dramatic increase in academic output. Journals, rather than being displaced by this flood, will become more powerful as filtering institutions. Established scholars will capture a disproportionate share of the gains. The result is a sharper stratification of the publication market, with consequences for who gets read and whose work shapes the field.

2. The Output Explosion

Academic writing has always been expensive. Not in money, but in time and cognitive effort. A credible literature review, a coherent theoretical framework, a

section connecting evidence to argument, each of these tasks consumes weeks of sustained effort. Generative AI compresses that timeline substantially. The marginal cost of producing a draft has fallen. This is not a prediction. It is an observation already visible in submission volumes. AI publications have experienced near-exponential growth, with AI conferences such as NeurIPS seeing surges in submissions in 2023 and 2024, and a large-scale PNAS study finds strong evidence of a sharp increase in AI-assisted academic writing across disciplines since 2023, with no significant difference between journals with and without AI policies.[1]

The cost reduction is real but uneven. AI assists with synthesis, structure, prose, even with generating ideas. Today, the papers that AI helps most are those that were already close to the frontier. Tomorrow, the AI will assist the entire scientific process. A large-scale Science study covering about 2.1 million preprints across arXiv, bioRxiv, and SSRN finds that LLM adoption is associated with substantial increases in researchers' manuscript output.[2] The result is a supply shock concentrated among already-active researchers. Output volumes rise. But the distribution of who produces that output is not uniform. Researchers who already publish frequently are best positioned to integrate AI into their workflow, because they have established pipelines, co-author networks, and institutional support. The marginal paper is easier to produce, and those who were already producing at scale benefit most from that reduction in marginal cost.

3. Journals as Attention Allocators

A common prediction holds that AI-assisted publishing will accelerate the decline of traditional journals.[3] The reasoning is intuitive. If anyone can produce a polished paper, the scarcity that journals once managed disappears, and alternative dissemination channels, preprint servers, substack-style platforms, curated feeds, fill the void.

This prediction mistakes the function of journals. Journals do not primarily produce papers. They filter them. Their value to readers is not access but selection. In an environment of low scarcity, selection becomes more valuable, not less. When

[1] Keigo Kusumegi et al., "Scientific Production in the Era of Large Language Models," Science 390, no. 6779 (2025): 1240–1243, https://doi.org/10.1126/science.adw3000; Yongyuan He and Yi Bu, "Academic Journals' AI Policies Fail to Curb the Surge in AI-Assisted Academic Writing," Proceedings of the National Academy of Sciences 123, no. 9 (2026): e2526734123, https://doi.org/10.1073/pnas.2526734123.

[2] Keigo Kusumegi et al., "Scientific Production in the Era of Large Language Models," Science 390, no. 6779 (2025): 1240–1243, https://doi.org/10.1126/science.adw3000.

[3] See, e.g., Gemma Conroy, "How ChatGPT and Other AI Tools Could Disrupt Scientific Publishing," Nature 622, no. 7982 (October 12, 2023): 234–236, https://doi.org/10.1038/d41586-023-03144-w.

every researcher produces twice as many papers, the reader's problem is not finding content. It is identifying what deserves attention.

The economics here follow the logic of attention markets. Herbert Simon observed that a wealth of information creates a poverty of attention.[4] Journals are attention brokers. Their role is to certify that a paper is worth reading before a reader invests the time to read it. As the volume of uncertified content increases, the premium on certification rises. Journals that maintain credible selection standards will attract more reader attention, not less. The expected equilibrium is not journal death but journal centrality.

This has a structural implication. The journals that function as credible filters will be more influential in shaping which ideas get absorbed into the field. Editorial choices will carry more weight. Reviewer time will be more scarce relative to submission volume. The gap between accepted and rejected will widen in practical terms, because the volume of rejected work will grow faster than the volume of accepted work. A journal accepting five percent of submissions when ten thousand papers are submitted per year is making a different kind of selection decision than one accepting five percent of two thousand. The filter tightens not in percentage terms but in absolute terms of what it excludes.

4. The Matthew Effect, Amplified

Robert Merton introduced the Matthew effect into the sociology of science in 1968, drawing on the Gospel of Matthew: to those who have, more will be given.[5] The observation was that cumulative advantage operates in academic careers. Early recognition generates citations, citations generate visibility, visibility generates invitations, and the cycle compounds. AI does not introduce this mechanism. It amplifies it.

The amplification works through two channels. The first is quality-correlated productivity. Researchers with strong intellectual foundations use AI to produce more work, and that work is better because the foundation is stronger. AI accelerates the steps that were bottlenecks, but it cannot supply what it does not receive. A researcher who enters the process with a sharp research question, a command of the relevant literature, and a clear theoretical position uses AI to move faster. A researcher without those inputs uses the same tools to generate volume. The gap between work that advances a field and work that merely adds to it may widen, because AI is a multiplier and what it multiplies is the intellectual capital brought to

[4] Herbert A. Simon, "Designing Organizations for an Information-Rich World," in Martin Greenberger, ed., Computers, Communications, and the Public Interest (Baltimore: Johns Hopkins Press, 1971), 40–41.

[5] Robert K. Merton, "The Matthew Effect in Science," Science 159, no. 3810 (January 5, 1968): 56–63, https://doi.org/10.1126/science.159.3810.56.

the task. That capital is not a function of career stage or institutional rank. It is a function of how clearly a researcher thinks before the first prompt is written.

The second channel is reputational. Journals make acceptance decisions under uncertainty. A submission from a lesser-known author requires editors and reviewers to evaluate the paper on its merits alone, with no prior signal of quality. A submission from a scholar with a sustained citation record, a Nobel Prize, or a recognizable institutional affiliation arrives pre-certified. The editor's prior probability that the paper is worth accepting is higher before anyone reads a word. This is not irrational behavior on the part of editors. It is Bayesian updating in conditions where reviewer time is scarce. When submission volumes rise, that scarcity increases, and the weight placed on reputational priors increases with it.

The mechanism economists describe as preferential attachment, in which nodes with high connectivity attract new connections at a higher rate, applies here directly. High-reputation authors attract editorial attention. Editorial attention produces acceptances. Acceptances produce citations. Citations reinforce reputation. The network feeds itself. AI, by increasing the total number of submissions competing for a fixed supply of top-journal slots, makes preferential attachment more consequential. The rich get richer not because the system becomes more biased, but because the competition intensifies around a constant bottleneck.

5. Stratification

The combination of output explosion and Matthew effect amplification produces a predictable structural outcome. The journal market stratifies more sharply.

Top journals increasingly publish work by established names. Empirical work already confirms the pattern. A study of political science journals found that editors favor authors affiliated with their home institutions,[6] and recent preprints testing LLMs as peer reviewers found that three out of four models exhibited bias toward well-known authors and all four favored prominent institutions.[7] This reflects the behavior of editors operating under greater selection pressure. It also tends to reinforce citation-based metrics such as the journal's h-index, as publishing work by already well-cited authors increases the likelihood of subsequent citations. But it has a distributional consequence. If the top journals fill with high-reputation authors

[6] Yaniv Reingewertz and Carmela Lutmar, "Academic In-Group Bias: An Empirical Examination of the Link between Author and Journal Affiliation," Journal of Informetrics 12, no. 1 (2018): 74–86, https://doi.org/10.1016/j.joi.2017.11.006.

[7] Rui Ye et al., "Are We There Yet? Revealing the Risks of Utilizing Large Language Models in Scholarly Peer Review," arXiv:2412.01708 (December 2, 2024), https://arxiv.org/abs/2412.01708; Pat Pataranutaporn et al., "Prestige over Merit: An Adapted Audit of LLM Bias in Peer Review," arXiv:2509.15122 (September 18, 2025), https://arxiv.org/abs/2509.15122.

publishing at higher volume than before, the available slots for lesser-known researchers shrink in absolute terms. A journal publishing one hundred articles per year that previously allocated thirty slots to emerging scholars may, under AI-driven submission pressure, find that established scholars now submit enough material to fill most of those slots with work that clears the quality bar.

Part of the response is likely to involve the expansion of publication outlets. This dynamic is already visible. New outlets appear, often open-access, often with faster review cycles. Gold open-access articles grew from 2% of global publications in 2003 to 30% in 2022,[8] and Springer Nature alone launched 68 new journals in 2024.[9] Many of these outlets serve researchers who do not have access to the networks that feed the top journals. These journals are not without value. But they occupy a lower position in the citation hierarchy, which means work published in them has less influence on the field's direction. The stratification of journals becomes the stratification of influence.

The long-run consequence is a publication market that looks less like a single competitive field and more like a set of parallel tracks. Track one: high-reputation authors publishing in high-visibility journals, producing the ideas that get cited, taught, and applied. Track two: a large volume of work circulating in secondary outlets, contributing to the literature in technical terms but rarely crossing into track one's citation networks. Emerging scholars, researchers outside elite institutions, and scholars from regions historically underrepresented in top publications will disproportionately find themselves on track two.

This is not a new problem. It is an old problem with a new accelerant.

6. Conclusion

Generative AI is accelerating the existing mechanisms of concentration in academic publishing. More output means fiercer competition for the slots that carry reputational weight. Fiercer competition advantages those who already have reputational weight. The journals that do credible filtering will become more central to how the field organizes its attention. The journals that do not will multiply and fragment.

The practical question for researchers, particularly junior ones, is not how to produce more. It is how to be read. In a market saturated with AI-assisted output, the returns to genuine intellectual originality, to asking questions that are actually new,

[8] National Science Foundation, National Center for Science and Engineering Statistics, "Open-Access Publishing in a Global Context," NSF InfoBrief NSF 25-347 (2025), https://ncses.nsf.gov/pubs/nsf25347; underlying data from National Science Board, Science and Engineering Indicators 2024, NSB-2023-33 (Alexandria, VA: National Science Foundation, 2023).

[9] Springer Nature, Open Access Report 2024 (2025), https://stories.springernature.com/oa-report-2024.

will be higher than they have been in decades. There is another escape route, and it runs through prose itself. Few readers would ask AI to summarize Kerouac instead of reading him. Few would replace the experience of reading Georges Perec with a bullet-point digest. The reading is the experience. Academic writing that achieves genuine literary quality, that makes the reader want to stay in the sentence, occupies a similar position. When the baseline of competent, AI-polished prose rises, the premium on writing that is not merely competent but distinctive rises with it. Scholars who can create (with AI or not) a real reading experience may escape the dynamics described above, because their work is consumed for what it is, not merely for what it concludes.

Thibault Schrepel
Vrije Universiteit Amsterdam