
The Economics of AI Slop: How Cost-per-Paper Alters the Academic Publishing Ecosystem

Rachel So

rachel.so@open.science

Abstract

The marginal cost of generating a research paper with large language models (LLMs) has fallen sharply, from thousands of dollars in researcher time to a few dollars of compute. This paper analyzes the economic consequences of that cost reduction for academic publishing. We argue that cheap AI-assisted paper generation does not merely accelerate scholarship; it reshapes incentive structures in ways that favor quantity over quality, amplify existing pathologies such as paper mills and predatory journals, and impose growing costs on the peer-review system. We introduce the concept of *AI slop* in the academic context: content that is superficially competent but lacks the original intellectual contribution expected of peer-reviewed research. Drawing on economic models of information markets, publishing incentives, and principal-agent theory, we characterize the equilibria that emerge when the cost-per-paper collapses. We show that without structural countermeasures, the publishing ecosystem faces a lemons problem in which low-cost, low-value papers crowd out high-cost, high-value ones. We conclude with policy recommendations for journals, funders, and academic institutions.

1 Introduction

Academic publishing rests on a tacit bargain: authors invest substantial time and expertise in producing original research, and journals certify that work through peer review. The economics of this system have always been shaped by the marginal cost of producing a paper. Historically, the dominant costs were human: months of experimental work, data analysis, and scholarly writing. These high entry costs acted as a natural filter, discouraging low-effort submissions.

Large language models (LLMs) disrupt this cost structure. A frontier LLM can draft a coherent scientific manuscript from a rough outline in minutes at a cost of under \$15 per complete paper [11]. This cost will continue to fall as models improve and inference prices decrease. The question is not whether AI will be used in academic writing, but how the resulting cost shift changes the strategic behavior of authors, publishers, and reviewers.

This paper studies the economic consequences of cheap AI-assisted paper generation. We use the term *AI slop* to describe AI-generated content that is superficially coherent but lacks genuine intellectual contribution. The term has been adopted in the broader discourse on AI-generated content to describe material produced with asymmetric effort and superficial competence [9, 8]. In the academic context, AI slop is research text that passes surface-level quality checks but does not advance knowledge.

Our central claim is that the cost reduction from LLMs is not simply a productivity gain. It is a structural shock that changes who submits papers, what they submit, and how journals and reviewers respond. The shock interacts with pre-existing pressures in academic publishing, including the “publish or perish” culture [6, 18], the growth of predatory journals [17, 16], and organized paper mills [4, 5].

The paper is organized as follows. Section 2 reviews background on publishing incentives and the cost structure of paper production. Section 3 characterizes the cost shock from LLMs and its immediate effects. Section 4 analyzes the market equilibria that emerge. Section 5 discusses the problem of detecting AI slop. Section 6 presents policy responses, and Section 8 concludes.

2 Background

2.1 Publishing Incentives and Publish-or-Perish

Academic career advancement in most institutions is tightly coupled to publication count and citation metrics [18]. This creates strong incentives to maximize the number of publications, sometimes at the expense of quality. Grimes et al. model how publication pressure interacts with false positive rates to undermine science trustworthiness, finding that decreasing funding exacerbates perverse incentive effects [6].

The publish-or-perish dynamic predates LLMs. Its consequences include salami-slicing of results across multiple papers, strategic journal targeting driven by metrics rather than audience fit, and increased rates of research misconduct. LLMs do not create these incentives, but they lower the marginal cost of acting on them.

2.2 Predatory Journals and Paper Mills

The open access movement, while broadly beneficial, introduced the article processing charge (APC) model in which publishers earn revenue per accepted paper. This creates a well-documented perverse incentive: low-status journals benefit financially from maximizing acceptance rates [17, 19]. Predatory journals exploit this incentive by charging fees while providing minimal or no real peer review [16, 12].

Paper mills represent a more systematic form of publishing fraud. These are organizations that produce fake or fabricated research papers and sell authorship to researchers who need publications for career advancement [4]. Analysis of peer-review data reveals that paper mills often operate by creating fake reviewer accounts and submitting fabricated reviews [5]. Paper mills already exploit low-cost document production; LLMs dramatically reduce their operating costs while making their output harder to detect.

2.3 Cost Structure of Paper Production

Before LLMs, the cost of producing an academic paper was dominated by researcher time. Typical estimates place the cost of a single experimental paper in the range of tens of thousands of dollars when researcher salaries and overhead are included. Even for purely computational or theoretical papers, skilled writing time is substantial. These costs served as an implicit quality filter: authors with nothing genuine to report had limited incentive to invest the effort required.

LLMs collapse the writing cost component to near zero. The AI Scientist system, for example, produces a complete machine-learning paper including literature review, experiments, and write-up for less than \$15 [11]. Even basic use of commercial LLMs to generate paper text costs only a few dollars per paper [3]. The research cost (experiments, data collection) remains high, but the writing and assembly cost does not.

3 The Cost Shock and Its Immediate Effects

3.1 What the Cost Reduction Changes

Let c_H denote the human cost of producing a paper of type H (high-effort, original research) and c_L denote the cost of producing a paper of type L (low-effort, AI-generated, little original contribution). Before LLMs, both costs were substantial. The ratio c_L/c_H was typically close to 1 for text-heavy fields: assembling and writing a paper required substantial effort regardless of the underlying research.

With LLMs, c_L falls dramatically. The ratio c_L/c_H can approach 0.01 or less when the primary bottleneck shifts from writing to original experimentation. This is not merely a cost reduction; it is a qualitative change in who finds paper production worthwhile.

Formally, an author with research output of quality q will submit if:

$$B(q) - c \geq 0, \tag{1}$$

where $B(q)$ is the expected career benefit from a publication of quality q and c is the production cost. When $c_L \approx 0$, even authors with $q \approx 0$ (i.e., minimal genuine contribution) may find submission rational if $B(0) > 0$. In most academic systems, $B(0) > 0$: any indexed publication provides some career benefit, and reviewers cannot perfectly observe q .

3.2 Observed Changes in Academic Writing

Empirical evidence already documents the impact of LLMs on academic text. Lin et al. analyzed 2.8 million abstracts from OpenAlex between 2020 and 2024 and found that ChatGPT introduction significantly increased lexical complexity in papers by non-native English speakers [10]. This finding illustrates a genuine benefit of LLMs: reducing linguistic barriers for researchers worldwide. However, the same mechanism that democratizes access also enables mass production of superficially sophisticated text.

Perkins et al. found that both automated detection tools and experienced academic staff struggle to reliably identify AI-generated content [13]. This asymmetric detectability is central to the economic problem: if journals cannot cheaply distinguish AI slop from genuine research, the signal value of publication declines.

3.3 The Shift from Quality to Volume

The fall in c_L shifts the optimal strategy for certain types of academic actors. Consider a researcher facing a fixed career evaluation period. If publication count matters and quality is hard to verify ex ante, then producing many low-cost papers dominates producing fewer high-cost ones whenever $B(\hat{q}) \approx B(q^*)$, where \hat{q} is the quality of an AI-assisted paper and q^* is the quality of a fully original one. Even a modest correlation between publication count and evaluation outcomes makes the volume strategy rational.

This logic is not new. Publish-or-perish pressure has long incentivized quantity [6]. LLMs simply lower the cost of acting on this incentive from a costly strategy to a nearly free one.

4 Market Equilibria Under Low Paper Production Costs

4.1 A Lemons Problem in Academic Publishing

Akerlof’s market for lemons [1] describes a market where quality is unobservable to buyers, leading high-quality sellers to exit and low-quality sellers to dominate. Academic publishing faces an analogous dynamic.

In the current market, journals are the “buyers” of papers, and quality certification (acceptance) confers value on authors. Readers and citing researchers are ultimate consumers of published knowledge. If the production cost of low-quality papers falls toward zero while detection costs remain high, the following equilibrium emerges:

1. The supply of low-quality (AI slop) submissions increases substantially.
2. Journals that cannot effectively screen face either increased acceptance of low-quality work or unsustainable reviewer burden.
3. High-quality venues tighten acceptance criteria, but the resulting rejection of genuine work as collateral damage reduces their appeal to high-quality authors.
4. Predatory and low-threshold journals absorb the overflow, growing in volume and apparent legitimacy through sheer publication count.
5. Citation networks increasingly include low-quality work, reducing the signal value of citations overall.

This is a standard adverse selection cascade. The key feature that makes LLMs novel is the speed at which the cascade can occur: a single coordinated actor with access to frontier LLMs can generate hundreds of superficially plausible papers per day at trivial cost.

4.2 Principal-Agent Problems in Peer Review

Peer review is a principal-agent relationship: journals (principals) delegate quality assessment to reviewers (agents) who have private information about paper quality but bear the cost of review effort. Under the existing system, reviewer effort is uncompensated, creating incentives for superficial review.

LLMs amplify this problem on two dimensions. First, the volume of submissions increases, raising the cost of thorough review. Second, some reviewers themselves use LLMs to generate reviews, further degrading quality assessment. Yu et al. demonstrate that existing AI text detection methods fail to reliably identify LLM-generated peer reviews, creating a second-order AI slop problem [21].

Saig et al. analyze the design of contracts for incentivizing quality in AI-assisted text generation [15]. They show that pay-per-token pricing creates a moral hazard: agents can substitute cheap models for expensive ones without detection. The academic equivalent is the substitution of AI-generated prose for genuine scholarly effort. Their result on cost-robust contracts suggests that quality incentives can be designed without knowledge of internal generation costs, a relevant insight for journal policy design.

4.3 Scale Effects and Platform Dynamics

The economic impact of AI slop is non-linear in volume. Below a threshold, journals can manually screen suspicious submissions. Above it, the screening cost becomes prohibitive, forcing either algorithmic detection (with its own false positive problems) or acceptance of degraded quality.

The publishing ecosystem also has platform characteristics. High-reputation venues attract high-quality authors partly because of their exclusivity. If AI slop floods submissions and forces tighter screening, the probability of genuine papers being rejected rises. This creates a reputational externality: each piece of AI slop submitted to a high-quality venue imposes a cost on all legitimate submitters by consuming reviewer time and potentially triggering false-positive screening errors.

4.4 Predatory Journals as Equilibrium Beneficiaries

Predatory journals are equilibrium winners in the low- c_L regime. Their value proposition, charging fees for guaranteed publication without real review, becomes more attractive as legitimate venues tighten screening. Researchers who produce AI slop but need indexed publications find predatory journals the natural outlet.

The APC structure of predatory journals already created misaligned incentives [17]. LLMs increase the supply of papers that need such outlets. Shamseer et al. document that predatory journals charge substantially lower APCs than legitimate open-access journals (median \$100 vs. \$1865), making them accessible to volume producers [16]. Combining cheap paper production with cheap publication creates a closed economic loop that requires minimal investment per publication credit.

5 Detection and Its Limits

5.1 The Detection Arms Race

AI text detection tools have proliferated alongside LLMs. Alhijawi et al. report accuracy improvements of up to 37.4% over baseline methods for detecting LLM-generated scientific text [2]. However, detection accuracy and false positive rates remain in tension: systems that flag most AI slop also incorrectly flag legitimate human writing, creating liability for journals that act on detection results.

Perkins et al. found that academic staff identified only 54.5% of AI-generated submissions as suspicious, and detection tool coverage was 54.8% [13]. These figures are from 2023 and apply to unobfuscated AI output. Advanced prompting techniques can substantially lower detection rates,

and researchers using LLMs for legitimate assistance produce text that overlaps with pure AI generation. The detection problem is not simply a technology challenge; it is an adversarial game in which detection and evasion co-evolve.

5.2 Quality Signals That AI Cannot Easily Fake

Not all dimensions of paper quality are equally susceptible to AI slop. The following signals are harder for current LLMs to fabricate:

- **Novel experimental results:** Data from original experiments require physical or computational resources that LLMs do not provide.
- **Reproducibility artifacts:** Code, datasets, and detailed protocols that reviewers can verify independently.
- **Longitudinal coherence:** A research program with consistent methodology and building results across multiple papers is harder to fabricate than isolated papers.
- **Community engagement:** Interaction in workshops, responses to reviewer comments, and collaborative work are signals of genuine participation.

These signals suggest a direction for structural reforms: shift evaluation weight from published text toward verifiable artifacts.

5.3 Hallucination as a Detection Signal

LLMs hallucinate: they generate plausible-sounding but factually incorrect content, including fabricated citations [14]. In the academic context, hallucinated references are detectable by automated bibliographic verification. Several journals have begun requiring that all cited papers be verified as real, a low-cost screening step that identifies a class of AI-generated submissions. However, authors can prompt LLMs to include only real citations, so this signal degrades as awareness of the check spreads.

5.4 AI Scientists and Genuine Research

It is important to distinguish AI slop from the emerging category of fully autonomous AI research systems. Lu et al. describe the AI Scientist, a system that generates research ideas, writes code, executes experiments, and writes complete papers [11]. Zhu et al. evaluate AI scientist systems critically, arguing that their fundamental bottleneck is execution capability rather than writing [22]. Hosseini et al. discuss the institutional risks of AI agents in research, including responsibility gaps and deskilling [7].

These systems are conceptually different from AI slop. An autonomous AI scientist that runs genuine experiments and produces verifiable results is contributing original knowledge, regardless of whether a human authored the prose. The economic problem we analyze arises not from AI that does research, but from AI that generates the appearance of research without the underlying substance. The distinction between real AI-assisted research and AI slop is crucial for policy design: interventions that target AI text will penalize genuine AI-assisted research alongside fraud.

6 Policy Responses

6.1 For Journals and Publishers

Artifact requirements. Journals should require submission of reproducibility artifacts (code, data, protocols) as a condition of review for computational and empirical papers. These artifacts shift part of the cost of evaluation from reviewers to automated verification tools and make it substantially harder to produce purely AI-generated submissions.

Disclosure and transparency. Mandatory disclosure of AI tool use in paper preparation, already adopted by many journals, raises accountability without banning legitimate AI assistance. Disclosure requirements do not solve the detection problem, but they create a paper trail that supports post-publication audits.

Reviewer compensation and load management. The reviewer burden imposed by increased submission volume is a real economic cost. Journals should consider tiered review processes, where a lightweight first-stage filter (automated plus handling editor) screens for obvious AI slop before papers reach human reviewers, protecting reviewer time.

Dynamic APC pricing. Journals using the APC model should consider pricing structures that disincentivize volume production. Graduated fees, discount caps, or institutional submission limits could reduce the economic attractiveness of mass submission strategies.

6.2 For Funders and Institutions

Evaluation metric reform. The publish-or-perish incentive is the root demand driver for AI slop. Funders and institutions that replace pure publication counts with quality-weighted metrics, including citation impact, artifact availability, and reproducibility scores, reduce the benefit side of the low-effort submission calculation.

Research integrity auditing. Funders could require statistical auditing of publication patterns as a condition of grant renewal. Unusually rapid publication rates, implausible co-authorship networks, or systematic bibliographic errors are detectable signals of AI slop production at scale.

Support for detection infrastructure. Investment in shared, open-source detection infrastructure benefits the entire publishing ecosystem. Individual journals lack the resources and data to train effective detectors; a consortium approach could provide better signal with lower false positive rates.

6.3 For Authors and Research Communities

Open science norms. Preregistration, open data, and open code requirements make it harder to substitute AI text for genuine research. These norms already exist in many fields and should be extended where possible.

Community-level standards. Research communities can establish norms around what AI assistance is acceptable. The key distinction is between AI that helps researchers express genuine ideas (acceptable) and AI that substitutes for the ideas themselves (unacceptable). Clear community standards give journals and institutions a baseline from which to enforce policies.

7 Discussion

The economic analysis presented here has several limitations. Our model of author behavior treats paper quality and production cost as the primary variables, abstracting from disciplinary differences, cultural contexts, and the heterogeneous nature of what constitutes an “original contribution.” The equilibria we describe are tendencies rather than deterministic predictions; the actual trajectory will depend on how quickly journals, funders, and detection tools respond to the cost shock.

There is also a genuine benefit to AI assistance in research writing that our analysis should not obscure. LLMs reduce barriers for non-native English speakers [10], help researchers communicate more clearly, and can accelerate the assembly of literature reviews. Yang and Zhang analyze how AI availability affects content production decisions, showing that the direction of incentive effects depends on the interplay between AI quality, copyright protection, and market structure [20]. A policy that successfully eliminates AI slop while also eliminating legitimate AI assistance would impose real costs on the research community.

The key challenge for policy design is separating the writing tool from the research contribution. A paper that presents genuine novel experiments but uses an LLM to improve prose quality is not AI slop. A paper that uses an LLM to fabricate the appearance of experiments is. Policies should target the absence of genuine contribution, not the presence of AI tools.

The parallel with earlier disruptions is instructive. The introduction of desktop publishing in the 1980s lowered the cost of producing professional-looking documents and contributed to the growth

of predatory conference proceedings. The rise of the internet lowered distribution costs and enabled the early growth of predatory journals. Each disruption created a temporary advantage for low-quality producers before the ecosystem adapted. LLMs represent a quantitatively larger shock because they lower not just distribution or formatting costs but the cost of generating persuasive scientific prose itself. The adaptation will need to be correspondingly more fundamental.

8 Conclusion

The marginal cost of producing a research paper has fallen sharply with the availability of capable LLMs. We have argued that this cost reduction is a structural shock to academic publishing with predictable adverse selection consequences. When the cost of generating superficially competent text approaches zero, the equilibrium involves more AI slop, more pressure on peer review, more growth in predatory publishing, and a degradation of the information value of publication credits.

These consequences are not inevitable. They are the outcome of a cost structure interacting with existing incentive systems. Targeted interventions, including artifact requirements, evaluation metric reform, disclosure norms, and detection infrastructure investment, can shift the equilibrium toward one in which the genuine benefits of AI-assisted research are captured while the production of AI slop remains economically unattractive.

The academic community faces a time-sensitive coordination problem. The institutions that act early to establish clear standards and structural safeguards will protect the reliability of the knowledge they produce. Those that wait will find the problem significantly harder to address once AI slop has become normalized in their publication records.

Acknowledgments and Disclosure of Funding

Generative AI has been used to prepare this paper.

References

- [1] George A. Akerlof. The market for “lemons”: Quality uncertainty and the market mechanism. *The Quarterly Journal of Economics*, 84(3):488–500, 1970.
- [2] Bushra Alhijawi, Rawan Jarrar, Aseel AbuAlRub, and Arwa Bader. Deep learning detection method for large language models-generated scientific content. *Neural Computing and Applications*, 37:91 – 104, 2024.
- [3] Adhari Alzaabi, Amira ALAmri, Halima Albalushi, Ruqaya Aljabri, and A. AAIAbdulsalam. Chatgpt applications in academic research: A review of benefits, concerns, and recommendations. *bioRxiv*, 2023.
- [4] Dorothy V. M. Bishop and Anna Abalkina. Paper mills: a novel form of publishing malpractice affecting psychology. *Meta-Psychology*, 2023.
- [5] A. Day. Exploratory analysis of text duplication in peer-review reveals peer-review fraud and paper mills. *Scientometrics*, 127:5965 – 5987, 2022.
- [6] D. Grimes, C. Bauch, and J. Ioannidis. Modelling science trustworthiness under publish or perish pressure. *Royal Society Open Science*, 5, 2018.
- [7] Mohammad Hosseini, Maya Murad, and David B. Resnik. Benefits and risks of using ai agents in research. *The Hastings Center Report*, 56:13 – 17, 2026.
- [8] Eric M Jones, Jane D Newman, Boyun Kim, and E. Fogle. Ai-generated “slop” in online biomedical science educational videos: Mixed methods study of prevalence, characteristics, and hazards to learners and teachers. *JMIR Medical Education*, 11, 2025.
- [9] Cody Kommers, Eamon Duede, Julia Gordon, Ari Holtzman, Tess McNulty, Spencer Stewart, Lindsay Thomas, R. So, and Hoyt Long. Why slop matters. *ArXiv*, abs/2601.06060, 2025.

- [10] Dingkang Lin, Naixuan Zhao, Dan Tian, and Jiang Li. Chatgpt as linguistic equalizer? quantifying llm-driven lexical shifts in academic writing. *ArXiv*, abs/2504.12317, 2025.
- [11] Chris Lu, Cong Lu, R. T. Lange, J. Foerster, Jeff Clune, and David Ha. The ai scientist: Towards fully automated open-ended scientific discovery. *ArXiv*, abs/2408.06292, 2024.
- [12] A. Memon. How to respond to and what to do for papers published in predatory journals? *Science Editing*, 2018.
- [13] Mike Perkins, Jasper Roe, Darius Postma, James McGaughran, Don Hickerson British University Vietnam, Vietnam, James Cook University Singapore, and Singapore. Detection of gpt-4 generated text in higher education: Combining academic judgement and software to identify generative ai tool misuse. *Journal of Academic Ethics*, 22:89–113, 2023.
- [14] Vipula Rawte, A. Sheth, and Amitava Das. A survey of hallucination in large foundation models. *ArXiv*, abs/2309.05922, 2023.
- [15] Eden Saig, Ohad Einav, and Inbal Talgam-Cohen. Incentivizing quality text generation via statistical contracts. *ArXiv*, abs/2406.11118, 2024.
- [16] Larissa Shamseer, D. Moher, Onyi Maduekwe, Lucy Turner, V. Barbour, R. Burch, Jocelyn P Clark, J. Galipeau, J. Roberts, and B. Shea. Potential predatory and legitimate biomedical journals: can you tell the difference? a cross-sectional comparison. *BMC Medicine*, 15, 2017.
- [17] Kyle Siler. Demarcating spectrums of predatory publishing: Economic and institutional sources of academic legitimacy. *Journal of the Association for Information Science and Technology*, 71:1386 – 1401, 2018.
- [18] S. Singh. Commentary: Publish or perish – musings of a young faculty. *Indian Journal of Ophthalmology*, 69:3725 – 3726, 2021.
- [19] L. Vo, David Armany, S. Bariol, Sriskanthan Baskaranathan, Tania Hossack, David Ende, and H. Woo. Financial barriers in urology publishing: an analysis of legitimate and predatory journals. *Anz Journal of Surgery*, 95:744 – 748, 2025.
- [20] S. A. Yang and Angela Huyue Zhang. Generative ai and copyright: A dynamic perspective. *ArXiv*, abs/2402.17801, 2024.
- [21] Sungduk Yu, Man Luo, Avinash Madasu, Vasudev Lal, and Phillip Howard. Is your paper being reviewed by an llm? investigating ai text detectability in peer review. *ArXiv*, abs/2410.03019, 2024.
- [22] Minjun Zhu, Qiujie Xie, Yixuan Weng, Jian Wu, Zhen Lin, Linyi Yang, and Yue Zhang. Ai scientists fail without strong implementation capability. *ArXiv*, abs/2506.01372, 2025.