ISSN–2278-5655

AMIERJ Aarhat Multidisciplinary International Education Research Journal

Volume–XV, Special Issues– I(b) Jan – Feb, 2026

OPEN ACCESS

Original Research Article

THE AI FLOOD: CONGESTION COLLAPSE IN SCIENTIFIC PEER REVIEW

*** Stavan Wadekar**

Department of Computer Science, Model College Dombivli, India.

Abstract:

Scientific peer review functions as the primary quality control mechanism in academic publishing. For decades the system remained manageable partly because producing a full manuscript required significant human effort which naturally limited submission volume. Generative artificial intelligence has changed this constraint by reducing the time and cost required to generate academic style text and enabling large scale submission growth. This introduces a new infrastructure level risk where the critical threat is not only low quality writing but sustained overload of editorial and reviewer capacity.

This study frames AI-enabled submission flooding as a system stability problem and models the peer review workflow as a two stage queueing pipeline consisting of editorial screening and peer review. A stress testing approach is applied by increasing manuscript arrival rates above baseline levels and evaluating system utilization and backlog formation and expected review delay. The results show a tipping point behavior where moderate increases in submission volume can push the system beyond its stability boundary causing persistent backlog growth and rapid inflation of review timelines from weeks toward months and years. This collapse occurs when baseline peer review utilization is already high, such that a moderate proportional increase in submissions is sufficient to push reviewer demand beyond available service capacity.

The findings indicate that content-based AI detection tools alone cannot prevent collapse under high volume conditions because congestion emerges when arrivals exceed service capacity regardless of manuscript origin. Therefore long term resilience requires defenses that reduce adversarial scaling and restore the balance between attacker cost and defender capacity through stronger identity verification and submission throttling and proof-of-personhood style controls.

Introduction:

The currency of science is trust, and the mechanism that secures that trust is peer review. For decades, we assumed this system was robust simply because the barrier to entry was high. Producing a manuscript used to require months of physical lab work, data processing, and careful writing. That difficulty acted as a natural gatekeeper, limiting the number of papers that editors had to manage [1]. But we ignored a design flaw. The system is built entirely on finite human attention. Editors and reviewers manage this work alongside their own research and teaching. The infrastructure has a hard capacity limit [2].

Generative AI has effectively nullified this capacity constraint. The time required to produce a manuscript has dropped to near zero. We have moved from an era of scarcity to an era of infinite supply. While AI tools help honest researchers, they also provide an asymmetric advantage to bad actors. Now, anyone can flood the submission pipeline with plausible, academic-style text without engaging in genuine research [3].

This is not just a quality control issue; it functions as a resource depletion attack on the scientific record. Even if a paper is entirely fabricated, a human editor must still allocate time to open, screen and reject it. When

ISSN–2278-5655

Aarhat Multidisciplinary International Education Research Journal

Volume–XV, Special Issues– I(b) Jan – Feb, 2026

OPEN ACCESS Original Research Article

multiplied by thousands this creates a massive backlog. Legitimate PhD students and scientists are caught in this congestion, facing extended delays because the processing queue is clogged with low-effort submissions.

The problem is compounded by the low cost of digital identities. Most journals allow registration with only an email address. This vulnerability enables "**Sybil attacks**," where a single actor creates multiple fake accounts to manipulate the system. Combined with automated text generation, a single attacker can generate infinite demand while the supply of human reviewers remains inelastic.

Current solutions are failing because they focus on detection rather than capacity. The field is racing to build tools that flag AI-generated text. But **queueing theory** suggests that detection comes too late. If submissions arrive faster than reviewers can process them, the queue grows mathematically unbounded. The delay occurs regardless of whether the fake paper is eventually caught.

This paper treats the peer review crisis as an infrastructure problem. Specifically, this study asks: at what exact rate of submission growth does the current review system become critically unstable? We use Queueing Theory to model the breaking point of the editorial workflow. Our data indicates that even a modest, sustained rise in submissions creates delays that persist long after the surge ends. We argue that the solution requires balancing the cost equation: we must implement defenses that make it expensive for attackers to flood the system without punishing legitimate researchers.

Literature Review:

Peer review is the primary quality filter for science, acting as an expert-based gatekeeper before any decision to publish is made [4]. But the system has a flaw: it was never built for infinite scale. Multiple studies define peer review as a capacity-constrained queue. When submissions spike, the number of available editors and reviewers cannot expand fast enough to catch up which results in creating inevitable delays [4]. The human cost is staggering. Recent data suggests researchers now spend over **100 million hours** a year just reviewing papers, a massive investment of time to keep the system running [5].

A major bottleneck here is "reviewer fatigue" [4]. This term describes the growing struggle editors face when trying to get experts to say "yes." As the flood of manuscripts rises, experts decline more invitations, and finding replacements takes weeks, not days [4]. This proves that slow review times are not just about judging quality—they are a logistical failure.

Why is the volume so high? The answer lies in academic incentives. Current evaluation systems are publication-driven, rewarding researchers who publish often and fast [4], [6]. This creates a dangerous imbalance: the number of papers grows exponentially, but the number of human reviewers grows only linearly. The pipeline is permanently clogged [4].

It is not just legitimate science clogging the pipe. Journals are also battling a wave of high-volume, low-value submissions. Editors are forced to waste time screening manuscripts that have almost no scientific value [4]. This "hidden workload" is critical because every hour spent rejecting bad papers is an hour taken away from reviewing good ones.

Worse still are the "paper mills." These are industrial operations that manufacture fake science using templates and fabricated data [7]. Investigations reveal that paper mills have successfully industrialized fraud, flooding journals with thousands of bogus manuscripts [8]. This attacks the very reliability of the scientific record and violates clear ethical guidelines on authorship [7], [9].

Now, Generative AI has accelerated this threat. Large Language Models (LLMs) allow anyone to generate fluent, academic-style text instantly, dropping the cost of writing a paper to near zero [10]. In blind tests, AI-generated abstracts have frequently fooled human reviewers [11]. While AI is a useful tool for drafting, it allows bad actors to flood the system faster than humans can possibly react [10].

Publishers have tried to fight back with AI detection tools [10]. These tools look for statistical patterns like "burstiness" to flag machine text. But they are unreliable. Studies show they frequently flag non-native English speakers as AI, creating a serious fairness issue [12].

More importantly, detection misses the point. Even a perfect detector cannot stop a Denial-of-Service attack. If the number of submissions exceeds the system's service capacity, the queue will crash regardless of whether the text is real or fake [4]. This is a system-level overload problem, not just a content problem.

This looks exactly like a "Sybil attack" in computer security. As defined by Douceur, a Sybil attack happens when one actor forges many identities to overwhelm a network [13]. In academia, unverified emails make creating fake authors easy and cheap [13]. Combined with AI text generation, attackers can scale their volume indefinitely while defenders are stuck with limited human attention.

Queueing theory gives us the math to model this collapse. Classical rules like **Little's Law** prove that waiting times explode once arrival rates get close to the service rate [14], [15]. While engineers use these models for internet traffic, they are rarely applied to peer review. This study fills that gap. We treat AI-driven flooding as an "Academic DoS" attack and use queueing based models to find the breaking point [14]. In short, the literature shows peer review is already cracking under pressure [4], [7]. AI simply adds a low-cost way to break it completely [10]. Since detection is failing, we need a new approach: modeling capacity to survive the shock.

Research Methodology:

This study uses a quantitative model to analyze how the peer review system behaves when submission volumes spike. The goal is simple we want to find out if the pipeline stays stable when submissions rise and estimate the exact breaking point where the backlog becomes permanent. Because we focus on infrastructure capacity, we do not judge the academic quality of the papers. We only measure how the system performs under pressure.

1. Research Approach

Our approach relies on simulation-based analysis. We treat peer review as a service system that has a hard limit on human capacity. Using the framework from Bianchi, Squazzoni, et al. [16], we view the process as a flow where manuscripts arrive continuously and move through two steps: screening and review. We use queueing theory to map this mathematically and check for congestion when the load gets heavy.

2. Research Design

The design is essentially a stress test. First, we analyze the system at a normal baseline rate. Then, we test it under higher arrival rates that mimic AI-driven growth. The objective is to watch the system transition from a stable state where delays are okay to an unstable state where the queue grows forever. This matches Mrowinski's findings [17], which showed that even small jumps in arrival intensity can trigger sudden phase transitions in review time.

3. Model and Framework

We model the pipeline with two distinct stages.

Stage 1 **Editorial Screening**: Editors check manuscripts for scope and formatting.

Stage 2 **Peer Review**: Experts evaluate the papers that survive screening.

We chose this structure because almost all journals screen papers before sending them to reviewers. The output from Stage 1 becomes the input for Stage 2.

We expect Stage 2 to be the main bottleneck since reviewer time is the scarcest resource.

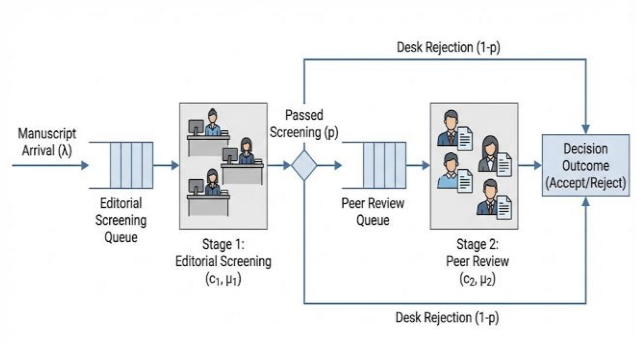

Figure 3.1 Peer Review Pipeline Model

(Description: Manuscript Arrival --> Editorial Screening Queue --> Peer Review Queue --> Decision Outcome)

4. Variables and Parameters

Here are the key variables in our model.

- ***λ* (Lambda)**: Manuscript arrival rate (submissions per unit time).
- ***μ1* (Mu 1)**: Service rate of editorial screening (manuscripts processed per editor per unit time).
- ***μ2* (Mu 2)**: Service rate of peer review (manuscripts reviewed per reviewer per unit time).
- ***c1***: Number of parallel editorial servers (editors available for screening).
- ***c2***: Number of parallel reviewer servers (active reviewers available for evaluation).
- ***p***: Fraction of manuscripts that pass editorial screening and proceed to peer review.

The effective arrival rate entering Stage 2 (Peer Review) is $p \times \lambda$.

5. Stability Conditions

The utilization (ρ) of each stage is defined as the ratio between the arrival rate and total service capacity.

Stage 1 Utilization:

$$\rho_1 = \frac{\lambda}{c_1 \mu_1}$$

Stage 2 Utilization:

$$\rho_2 = \frac{p\lambda}{c_2 \mu_2}$$

The system remains stable only when both stages satisfy the stability condition:

$\rho_1 < 1$ and $\rho_2 < 1$

The equality condition $\rho_2 = 1$ defines the system's stability threshold. At this point, the effective arrival rate into peer review equals total reviewer service capacity. Any sustained increase beyond this threshold results in mathematically unavoidable backlog growth and unbounded review delays.

When utilization approaches 1.0, the system becomes highly congested and waiting times increase sharply. When utilization exceeds 1.0, backlog growth becomes persistent and delays become unbounded in the long run.

6. Data and Assumptions

We use modeling assumptions because major publishers do not share their internal data. Our baselines are estimates that reflect a high-workload environment. For service rates, we rely on Kovanis et al. [18], who estimated the global reviewer pool and the time cost per paper. We assume manuscript arrivals follow a random **Poisson** process. We also assume service times are exponential. These are standard assumptions in queueing theory [14] that allow us to calculate stability mathematically.

Crucially, we assume that during a surge, submissions can rise fast but reviewer capacity stays flat. This reflects reality editors cannot magically find more free time just because more papers arrive.

7. Baseline Parameters

Symbol	Parameter Description
λ	Baseline submission arrival rate
p	Fraction passing screening
c_1	Number of editors for screening
μ_1	Editorial screening rate
c_2	Number of active reviewers
μ_2	Peer review service rate

Table 3.1 Baseline Model Parameters

ISSN–2278-5655

Aarhat Multidisciplinary International Education Research Journal

Volume–XV, Special Issues– I(b) Jan – Feb, 2026

OPEN ACCESS Original Research Article

We picked a baseline where the system is already busy. This reflects the real pressure many fields face today. The baseline configuration is chosen to reflect contemporary high load conditions in peer review, where peer review utilization is already close to saturation. In queueing terms, this corresponds to baseline values of $\rho 2$ that are significantly below but close to the stability threshold.

8. **Experiment Setup**

The experiment has two steps. First we define the Normal rate $\lambda 0$. Second we apply a Stress Factor to simulate an AI-enabled flood. The new rate is:

$\lambda stressed = (1+\alpha)\lambda 0$

We test stress levels using uniform 20% increments (20%, 40%, 60%, 80%, and 100%) relative to the baseline. This fixed-step stress design is used to systematically evaluate system behavior under progressively increasing submission pressure, focusing on qualitative stability transitions rather than precise empirical thresholds. The baseline configuration is chosen to reflect contemporary high-load conditions in peer review, where peer review utilization is already close to saturation. In queueing terms, this corresponds to baseline values of $\rho 2$ that are significantly below but close to the stability threshold.

9. Evaluation Metrics

We judge system health using four metrics.

- Utilization (ρ): Demonstrates how close the system is to maximum capacity.
- Backlog Size (Lq): Represents the number of manuscripts waiting in the queue.
- Expected Waiting Time (Wq): Represents the average delay experienced by a manuscript in the queue.
- Review Cycle Time: Represents the overall time taken from submission to final decision.

To link backlog and wait time, we use Little's Law [15].

$L = \lambda W$

These numbers help us spot the tipping point where manageable delays turn into a crisis.

1.0 Tools Used

We run this analysis using spreadsheets and standard simulation logic. We compute stability using the formulas above. The output includes tables and graphs showing how backlogs grow as stress rises. This method is fully reproducible in Excel

1.1 Study Limitations

This method has limits. We focus on capacity, not quality. We do not measure if the reviews are actually good. We assume average speeds even though some papers take longer than others. We also skip complex details like Revise and Resubmit loops. Finally, we treat volume as the main stressor, ignoring things like editor burnout. Despite this, the model gives us a clear framework to test peer review stability in the age of AI.

Findings and Results:

1 **Overview of Experimental Results**

This section presents the outcomes of the queueing based stress test performed on the peer review pipeline. The system was evaluated under a baseline submission rate and then under increased manuscript arrival rates to represent AI enabled submission growth.The objective was to observe the stability transition of the peer review system under rising submission pressure.

2. **Stress Test Results Table**

The quantitative results of the stress test are summarized below. The data reveals that the system maintains stability only within a narrow margin of submission growth.

ISSN–2278-5655
Aarhat Multidisciplinary International Education Research Journal
Volume–XV, Special Issues– I(b) Jan – Feb, 2026
OPEN ACCESS
Original Research Article

Submission Increase (%)	Capacity Regime	System Status	Backlog Behavior	Expected Review Delay
0% (Baseline)	Below capacity	Stable	Clears consistently	Weeks
+20%	Near capacity	Strained	Clears slowly	Weeks to Months
+40%	Above capacity	Unstable	Persistent growth	Months to Years
+60%	Far above capacity	Unstable	Rapid accumulation	Years
+80%	Severely overloaded	Collapsed	Explosive growth	Indefinite
+100%	System failure	Collapsed	Severe backlog	Unrecoverable

Table 4.1: Stress Test Outcomes Under Uniform 20% Submission Increments

3. Backlog Growth Under Stable and Unstable Conditions

Figure 4.1 illustrates backlog formation over time under three operating conditions. Under baseline load the backlog remains close to zero and clears consistently. Under near capacity conditions (+20%) the backlog rises slowly and becomes difficult to clear. Under unstable conditions (+40% and above) the backlog increases continuously because the arrival rate exceeds service capacity. This indicates that the peer review pipeline cannot recover once utilization crosses the stability boundary.

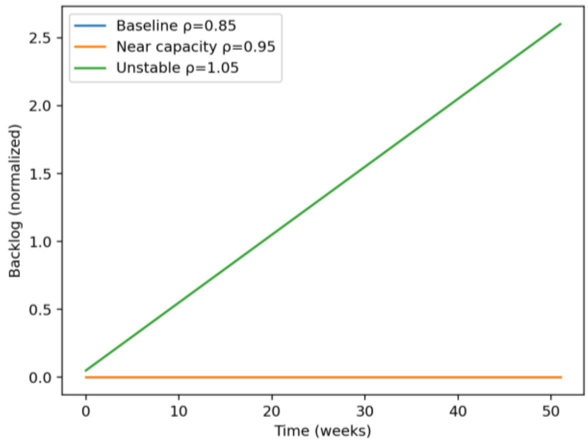

Figure 4.1: Backlog Growth Over Time (Baseline vs. Unstable)

4. Waiting Time Inflation Near the Stability Boundary

Figure 4.2 presents the expected trend of waiting time increases as submission volume rises. The results demonstrate that waiting time increases sharply as utilization approaches 1.0. When the system becomes unstable waiting time does not remain bounded and review timelines can expand drastically. This reflects a tipping point behavior where even moderate submission growth can produce disproportionately large delay increases.

ISSN–2278-5655

AMIERJ Aarhat Multidisciplinary International Education Research Journal

Volume–XV, Special Issues– I(b) Jan – Feb, 2026

OPEN ACCESS Original Research Article

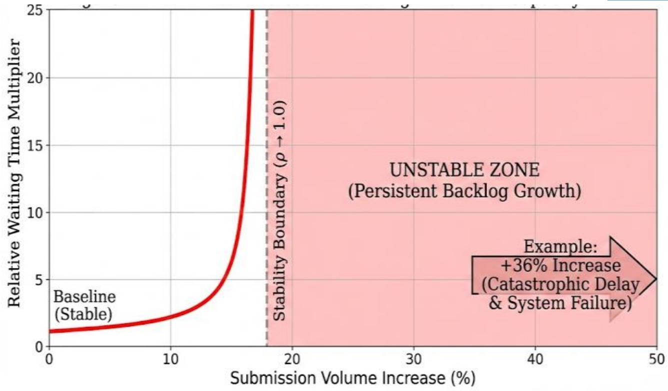

Figure 4.2: Waiting Time Inflation vs. Submission Increase

5. **Key Numerical Findings**

The stress test results highlight the following key findings:

Baseline Stability: Under baseline conditions the system remains stable and review delays remain within weeks.

Critical Threshold: With a 20% increase the system approaches critical capacity and delays begin to extend from weeks toward months.

Breaking Point: At a 40% increase, the system crosses the stability limit because peer review capacity is already near saturation at baseline, and the proportional increase in arrivals pushes utilization beyond unity, causing permanent backlog growth.For illustration, if baseline peer review utilization is $p2 = 0.72$, a 40% increase in submission volume yields $p2 = 1.4 * 0.72 = 1.01$, which exceeds the stability threshold and results in unavoidable backlog divergence.

System Collapse: At 60% and above the system becomes critically unstable and the accumulated backlog leads to long term delay escalation.

Cycle Expansion: Under high volume conditions review cycle time expands from weeks to months and can extend to years as congestion persists.

6. **Interpretation of Results**

These results indicate that the peer review pipeline is highly brittle under volume based stress. The infrastructure does not fail gradually. Instead it displays a sharp transition from stable performance to persistent overload once submission growth exceeds available human processing capacity. This confirms that under AI-enabled flooding the dominant failure mode is system level congestion rather than only content level authenticity.

7. **Sensitivity Check Based on Reviewer Capacity**

A brief sensitivity check suggests that increasing reviewer capacity can delay instability but may not eliminate collapse under sustained submission inflation. Even if reviewer availability improves slightly the system remains vulnerable because submission volume can scale faster than human review capacity. Therefore long term stability requires mechanisms that control the arrival rate rather than relying only on expanding service capacity.

8. **Empirical Validation: Case Study Analysis**

The theoretical thresholds identified in this study are reflected in the recent operational shifts of major scientific venues. For instance, the **NeurIPS** conference serves as a primary example of the "System Failure/Collapsed" regime. Between 2020 and 2025, submissions to the NeurIPS Main Track escalated from approximately 9,467 to 21,575 , representing a staggering 128% increase. This volume far exceeds the 100% stability boundary

ISSN–2278-5655

Aarhat Multidisciplinary International Education Research Journal

Volume–XV, Special Issues– I(b) Jan – Feb, 2026

OPEN ACCESS Original Research Article

modeled in Table 4.1. To maintain throughput, the conference was forced to recruit over 20,500 reviewers .However, this extreme expansion of the reviewer pool has been linked to significant review anomalies and controversial rejections due to venue capacity constraints. [23]

Similarly, **CVPR** (Computer Vision and Pattern Recognition) provides a clear case of a system operating within the "Strained" boundary. In 2024, the venue received 11,532 submissions , marking a 26% increase over the previous year. This empirical growth aligns closely with the +20% "Near Capacity" threshold predicted by the queueing model. In this regime, the model anticipates that backlogs clear slowly and require maximum resource utilization. This is validated by CVPR organizers, who noted that program committees are now expending "all their energies" simply to ensure papers find reviewers. [25]

Beyond specific conferences, broader multi-discipline trends confirm the predicted inflation of review delays. In fast evolving fields like Artificial Intelligence, the journey from submission to publication can now take six months to two years. This is evidenced by **PLOS ONE**, which reported a median publication time of 213 days in 2024. Such data confirms the model's prediction that as a system approaches its stability boundary, waiting times (Wq) do not remain bounded and expand drastically from weeks toward months and years. [22]

Discussion:

1. Peer Review as a Capacity-Limited Infrastructure

The findings of this study suggest that scientific peer review is best conceptualized as a capacity-limited infrastructure, rather than solely as a mechanism for content verification. Under typical operating conditions, the system functions close to its service limit because both editorial screening and peer review rely on scarce and inelastic human time. When manuscript arrivals remain below total processing capacity, backlogs are contained and decision timelines remain manageable. However, as arrival rates approach the system's capacity boundary, even marginal increases in submissions lead to disproportionately large increases in delay. This behavior demonstrates that peer review operates as a stability-sensitive pipeline, exhibiting sharp performance degradation near saturation.

The collapse observed at a 40% increase in submissions should not be interpreted as a universal threshold. Instead, it reflects baseline conditions in which peer review already operates near capacity. Under lower baseline utilization, substantially larger increases would be required to trigger instability, whereas under more saturated conditions, collapse would occur at lower levels of additional stress.

2. The Dominance of Congestion Over Quality

A major implication is that the dominant failure mode under AI-enabled submission growth is not limited to low-quality or AI-written manuscripts. The primary risk is congestion itself. Even if a large fraction of flooded manuscripts are eventually rejected, they still consume screening resources, reviewer recruitment effort, and decision processing time. As a result, the cost of overload is paid by the entire research community, including genuine authors whose manuscripts experience extended review cycles [19]. This explains why peer review delays are not simply a quality issue but a throughput and coordination issue.

3. Structural Limitations of AI Detection

The results also highlight why detection-based strategies are structurally insufficient as a long-term defense. AI detection tools attempt to classify text based on linguistic patterns and probabilistic signals. However, capacity collapse occurs when submission inflow exceeds service capability, regardless of whether the content is human-written or machine-generated. In other words, a system that is already near saturation cannot rely on partial filtering because even a small percentage of unchecked submissions can maintain overload conditions. Therefore, detection can reduce some fraudulent volume, but it cannot guarantee stability under sustained high arrival rates.

4. **The Asymmetry of AI Assistance**

An important counter-argument is that reviewers may also utilize AI to increase review speed and raise service capacity. While AI assistance can reduce some workload such as summarization, grammar correction and initial screening support but it does not remove the need for human responsibility. Reviewers remain accountable for judging novelty, methodological correctness and ethical concerns. In numerous fields, evaluation requires domain expertise, statistical understanding and careful verification that cannot be safely delegated to automated systems. This creates an asymmetry where attackers can scale submission volume rapidly with minimal cost, while defenders remain limited by human verification responsibility and reputational risk.

5. **The Sybil Attack Context**

The security framing of the problem further strengthens this interpretation. When submission systems rely on weak identity checks the cost of generating multiple author identities becomes low. This is comparable to a Sybil attack, where one actor multiplies identities to gain influence or overwhelm a system that assumes one identity corresponds to one real participant [20]. In academic publishing, this means that identity-based assumptions embedded in submission workflows can be exploited to create large-scale flooding. Under such conditions, the peer review pipeline resembles a Denial-of-Service (DoS) scenario where the attacker increases the arrival rate while the defender cannot expand service capacity at the same speed.

6. **Policy Implications and Resilience**

From a policy perspective, these findings suggest that improving peer review resilience requires mechanisms that reduce adversarial scaling and restore the balance between attacker cost and defender capacity. Measures such as stronger identity verification, submission rate limits per person, and institutional authentication can introduce friction and reduce the feasibility of mass identity creation. These controls do not need to eliminate AI usage entirely but they must prevent unlimited, low-cost flooding. Long-term solutions may require "proof-of-personhood" style controls that ensure each submitting identity is tied to one real human and that submission rights cannot be multiplied cheaply. Such mechanisms can preserve openness for legitimate researchers while preventing systematic overload that threatens the sustainability of peer review.

Conclusion:

By characterizing the peer review pipeline as a capacity constrained queueing system this study demonstrates its inherent structural vulnerability to AI amplified submission volumes. Simulations reveal that because qualified reviewer availability is inelastic even marginal increases in arrival rates when approaching service capacity precipitate exponential backlog accumulation and review latency. Crucially these findings establish that content-based detection is a futile defense against system saturation. The threat mechanism is volumetric overload not merely textual inauthenticity. Long term stability therefore demands a paradigm shift from content analysis to resource governance. To prevent systemic collapse publishing infrastructures must implement system level friction such as proof-of-personhood protocols and submission throttling and strong identity verification to re-impose economic costs on attackers and protect finite human attention reserves.

Limitations and Future Scope:

This study has specific limitations as it models peer review as a simplified queueing system focused on capacity and arrival dynamics. Consequently, the model does not evaluate manuscript quality, heterogeneity in reviewer expertise, or variations in review depth across disciplines. It also relies on average service rates, whereas real-world review completion times vary widely depending on manuscript complexity, reviewer workload, and editorial decision policies. In addition, the model does not explicitly represent resubmissions, revision rounds or withdrawal behavior which can further influence backlog patterns. Future work can refine this research by using field-specific parameter estimates collected from journals, conferences, and publisher reports and by extending the model to include multiple review rounds, desk rejection variability, and reviewer assignment constraints. Further studies should also evaluate hybrid defenses that combine identity-based throttling with risk-based screening and automated triage systems [21] to better protect peer review stability while preserving fairness and openness for genuine researchers

References:

1. *M. Ware and M. Mabe, "The STM Report: An overview of scientific and scholarly journal publishing," STM Association, The Hague, 2015.*
2. *J. Brainard, "New tools show promise for tackling paper mills," Science, vol. 380, no. 6645, pp. 568–569, 2023.*
3. *C. Stokel-Walker, "ChatGPT listed as author on research papers: many scientists disapprove," Nature, vol. 613, no. 7945, pp. 620–621, 2023.*
4. *J. H. Ellwanger and J. A. B. Chies, "We need to talk about peer review: Experienced reviewers are not endangered species, but they need motivation," Journal of Clinical Epidemiology, vol. 125, pp. 201–205, 2020.*
5. *B. Aczel, B. Szaszi, and A. O. Holcombe, "A billion-dollar donation: estimating the cost of researchers' time spent on peer review," Research Integrity and Peer Review, vol. 6, no. 1, pp. 1–8, 2021.*
6. *T. Tennant et al., "A multi-disciplinary perspective on emergent and future innovations in peer review," F1000Research, vol. 6, p. 1151, 2017.*
7. *J. Christopher, "The raw truth about paper mills," FEBS Letters, vol. 595, no. 13, pp. 1751–1756, 2021.*
8. *H. Else and R. Van Noorden, "The fight against fake-paper factories," Nature, vol. 591, no. 7851, pp. 516–519, 2021.*
9. *COPE Council, "COPE Guidelines: A Short Guide to Ethical Editing for New Editors," Committee on Publication Ethics, May 2019.*
10. *M. Khalifa and A. Albadawy, "AI in academic writing: A game changer or a threat to integrity?" Computer Methods and Programs in Biomedicine Update, vol. 5, p. 100145, 2024.*
11. *D. Gao et al., "Comparing scientific abstracts generated by ChatGPT to real abstracts with detectors and blinded human reviewers," npj Digital Medicine, vol. 6, no. 1, p. 75, 2023.*
12. *W. Liang et al., "GPT detectors are biased against non-native English writers," Patterns, vol. 4, no. 7, p. 100779, 2023.*
13. *J. R. Douceur, "The Sybil attack," in First International Workshop on Peer-to-Peer Systems (IPTPS), Springer, 2002, pp. 251–260.*
14. *L. Kleinrock, Queueing Systems, Volume 1: Theory. New York: Wiley-Interscience, 1975.*
15. *J. D. C. Little, "A proof for the queuing formula: $L = \lambda W$," Operations Research, vol. 9, no. 3, pp.*

383–387, 1961.

16. *F. Bianchi, F. Grimaldo, G. Bravo, and F. Squazzoni, "The peer review game: an agent-based model of scientists facing resource constraints," Scientometrics, vol. 116, no. 3, pp. 1401–1415, 2018.*
17. *M. J. Mrowinski, A. Fronczak, and P. Fronczak, "Review time in peer review: quantitative analysis and modelling of editorial workflows," Journal of Informetrics, vol. 10, no. 2, pp. 533–545, 2016.*
18. *M. Kovanis, R. Porcher, P. Ravaud, and L. Trinquart, "The global burden of journal peer review in the biomedical literature: Strong imbalance in the collective enterprise," PLOS ONE, vol. 11, no. 11, p. e0166387, 2016.*
19. *M. Breuning, L. Backstrom, J. Brannon, B. Gross, and M. Widmeier, "Reviewer fatigue? Why scholars decline to review their peers' work," PS: Political Science & Politics, vol. 48, no. 4, pp. 595–600, 2015.*
20. *M. Platt, D. Platt, and P. McBurney, "Sybil attack vulnerability trilemma," International Journal of Parallel, Emergent and Distributed Systems, vol. 39, no. 4, pp. 446–460, 2024.*
21. *A. Checco, L. Bracciale, P. Loreti, S. Pinfield, and G. Bianchi, "AI-assisted peer review," Humanities and Social Sciences Communications, vol. 8, 2021.*
22. *PLOS, "Explore research metrics at PLOS (2024 Metrics)," Sept 2025.*
23. *IntuitionLabs, "NeurIPS 2025: A Guide to Key Papers, Trends & Stats," Dec 2025.*
24. *H. Kataoka, "CVPR 2025 Report: Meta Insights into Trends and Tendencies," June 2025.*

Cite This Article:

Wadekar S. (2026). *The AI Flood: Congestion Collapse in Scientific Peer Review.* **In Aarhat Multidisciplinary International Education Research Journal**: Vol. XV (Number I, pp. 189–199)

Doi: https://doi.org/10.5281/zenodo.18638145