Episode Book

EPISODE #6470

Why Computer Science Publishes in Conferences

PUBLISHED	5 November 2026 at 20:58 UTC
DURATION	27:31
LISTEN	myweirdprompts.com/episode/computer-science-conference-proceedings-peer-review-13613/
HOSTS	Corn, Herman, Lars, Hilbert
SCRIPT MODEL	z-ai/glm-5.3
VOICES	chatterbox-regular (chatterbox-local)
PIPELINE	V5.26
SOURCES	25 cited by the research (listed at the end)
BOOK GENERATED	6 November 2026 at 11:12 UTC

SYNOPSIS

Computer science treats conference proceedings, rather than traditional journals, as its permanent archival record of research. This episode examines the mechanics behind academic publishing, the reality of acceptance rates, and the noisy nature of peer review revealed by experiments at major venues like NeurIPS. It explores why fields like speech recognition accept nearly half of submissions while machine learning caps admission at a quarter, and why hallway conversations remain the true venue for scientific progress.

TRANSCRIPT



Corn

Four thousand and thirty seven papers in one rejection cycle. I read that number this morning and needed a lie-down.



Herman

And that's just one conference. Welcome back, everybody. Corn, what's Daniel sent us?



Corn

Daniel's been down a rabbit hole. He's been fine-tuning ASR models for code-switching, which is the problem of recognizing speech when the speaker jumps between languages mid-sentence, and the deeper he goes, the more the literature turns out to live in conference proceedings. Interspeech, ICASSP, places like that. He's not an academic, doesn't want to be one, and he's asking a question that I suspect a lot of people outside the walls have quietly held for years: what are these things actually for? He quotes the rosy version, like-minded people sharing ideas for the greater good, and notes that anecdotal reality keeps intruding, acceptance-rate bragging, prestige games. His phrase is that he wants a circumspect overview, not a takedown. What does presenting a paper mean, why are proceedings an official record when we have every publishing tool ever invented, and what's really going on beyond the podium.



Herman

He also used the phrase meeting minutes written in quill, which I enjoyed, because it captures exactly what the proceedings sound like to an outsider.



Corn

So let's start with the obvious question. What are these things for, when you strip away the brochure language?



Herman

The reframe that helped me most, when I was reading around for this, is that the paper is the artifact and the conference is the process. Two sociologists, González-Santos and Dimond, wrote a review in twenty fifteen arguing that conferences are places where knowledge can be, their words, observed in the making. Not reported. Made. The hallway



Herman (continued)

argument that kills a bad assumption before it ever reaches print. That's the unit of value, and no PDF captures it.



Corn

Because the brochure version treats the talk as the product and the people as decoration. It's the other way around.



Herman

I think it mostly is. There's a formal layer, the papers and the proceedings and the acceptance decisions, and an informal layer, the corridors and the dinners and the reunion of people who only see each other at this one event every year. The information-science people have this lovely term, the invisible college, the network of researchers who exchange work before it's published. Once a year, the invisible college gets a body. It books a hotel.



Corn

The ghost gets a lanyard.



Herman

And a name tag. Sá, Ferreira and Serpa did a study in twenty nineteen on conference participation and their framing is that social networks are a co-equal function to dissemination. Not a side effect, not networking in the LinkedIn sense. A co-equal function of the event itself. So when Daniel hears that the idealistic vision is divorced from reality, my answer is that the vision isn't false, it's just described at the wrong layer. The ideas do get shared. They get shared at dinner, by forty people, and none of it is in the proceedings.



Corn

Meanwhile the formal layer, which is all outsiders ever see, is the part we're about to take apart.



Herman

And that gets us to the acceptance rate. First thing everyone points to as the measure of a conference's worth. So let's talk about that number, and what it actually measures.



Corn

I'll open with the flat observation, and you can put the color in. NeurIPS has held between roughly twenty-four and twenty-six percent for years. This past year, twenty-four point five two percent. The year before that, twenty-five point eight. Meanwhile submissions went from about nine and a half thousand in twenty twenty to twenty-one and a half thousand last year. More than doubled. The rate didn't move.



Herman

Which means the thing that grew is the rejection count. From around seven thousand rejections to over sixteen thousand in five years. So when a venue brags that it's selective, a big part of what it's saying is that a lot of people showed up. The bar isn't rising. The queue is.



Corn

That's the sentence I'd put on the brochure. NeurIPS: the queue is very long.



Herman

It's a fixed capacity decision wearing a quality costume. Nobody at NeurIPS sits down each year and says, the science got better, let's accept twenty-five percent instead of thirty. They hold the number because the number is the brand, and the submissions grow because the field grows and because the incentive structure pushes everyone to submit everything.



Corn

And the prestige hierarchy treats that number as if it measured something stable. Does it?



Herman

Here's the strongest piece of evidence anyone has ever produced, and it comes from the conference's own organizers. Twenty twenty-one, NeurIPS ran what they called a consistency experiment. They took ten percent of the submissions, over eight hundred papers, duplicated them, and had each copy reviewed by two completely independent committees. Neither committee knew the experiment was happening.



Corn

The same paper, two juries.



Herman

Two juries, full process, reviewers, area chairs, everything. Twenty-three percent of those papers got different answers. And here's the detail that should end acceptance-rate bragging at dinner parties: among papers accepted by one committee, half, a hair over fifty percent, were rejected by the other.



Corn

Half. So acceptance at one of the most prestigious venues in the field is a coin with a slight thumb on it.



Herman

The authors themselves computed the baseline. If acceptance were literally random at that year's rate, you'd get thirty-five percent inconsistency. The real figure was twenty-three percent. Their own words were that the outcome is closer to the random baseline than it is to zero. It's a thirty-five percent improvement over a coin flip, which sounds respectable until you remember it's being sold as the pinnacle of scientific vetting.



Corn

It gets worse the higher up you go, doesn't it. That's the part I find beautiful.



Herman

It's beautiful in the way a bridge collapsing in slow motion is beautiful. If you look at the most selective tier, the orals and spotlights, the committees together considered fifty-four duplicated papers for that tier. They agreed on three. Out of fifty-four. And of the spotlight papers one committee chose, more than half were outright rejected by the other committee. Spotlight designation, the thing that makes careers, is very close to a coin flip at the margin.



Corn

Thirteen of twenty-five. I could get those odds arguing with Herman about dinner.



Herman

You generally do. Now, to be fair to the people running these things, and Daniel asked for circumspect, not a takedown, the twenty twenty-five program chairs were remarkably candid about why the noise exists. Scale. They had twenty thousand reviewers. You cannot recruit twenty thousand qualified reviewers, so you get first-time area chairs, reviewers who aren't topically qualified for the paper they drew. And they openly described the discretionary layer: area chairs fighting for papers they liked even when reviewers disagreed, and often winning, or flagging problems after rebuttal and rejecting papers that had high ratings.



Corn

That's the politics beyond the podium, described by the people running the podium.



Herman

In a blog post, voluntarily. I find that almost more notable than the randomness. The institution knows.



Corn

Here's my corrective, though, and it's the one that matters for Daniel's field. Interspeech, which is the flagship venue for exactly the speech work he's been reading, accepted forty-nine point nine percent last year. Record-breaking number of papers, they said, celebrating it. ICASSP runs in the mid-forties. Double the selectivity gap of NeurIPS, and nobody in speech considers Interspeech disreputable. Quite the opposite.



Herman

It's the top of the field. And that tells you the prestige hierarchy attached to acceptance rate is a cultural fact, not a measured quality difference. Speech culture decided a fifty percent rate means we serve the community. Machine learning culture decided a twenty-five percent rate means we're elite. Same activity, different mythology. When I practiced medicine, the journal was the prize and the conference was the trailer for it. Computing picked the opposite convention, and I don't think one is more honest than the other.



Corn

So the answer to why conferences distinguish themselves by acceptance rate is: because the field decided to. The number measures the size of the queue and the capacity decision, and everything past that is branding.



Herman

One more wrinkle, because it cut me when I read it. Ninety-two papers at that NeurIPS were accepted only because the experiment ran, meaning at least one of their two copies cleared. Four were rejected because one committee found a fatal flaw the other missed. Both directions. The noise is symmetric. It isn't bias so much as weather.



Corn

Ninety-two people walked away with a line on their CV because a statistician got curious.



Herman

And they'll never know which ones they were.



Corn

Okay. So the sorting mechanism is noisy. That raises the bigger question. What is the thing being sorted? What does it actually mean to present a paper, and why do the proceedings still matter?



Herman

The proceedings question has a real answer, and it's the part of this whole subject I find least cynical. In computing, the conference isn't a trailer for the journal. The conference is the journal. The Computing Research Association says it flat out: for most of computer science, the flagship venue is a peer-reviewed conference and its published proceedings are the archival record. A CACM piece from two thousand eight put it that computing is the only scientific community that considers conference publication the primary means of publishing results. Goes back to the early nineteen eighties, because the discipline was young and fast-moving and journals took a year.



Corn

That actually makes sense to me. A field where a result is stale in eight months can't wait for a journal cycle. The conference is a deadline that fires on schedule whether the paper is ready or not.



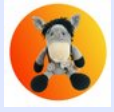
Herman

That's the double edge, Corn. The deadline firing on schedule whether the paper is ready or not is a fair description of how a lot of bad papers get in. But the archival function is real. What Daniel's quill-and-minutes intuition misses is that the value isn't the format. It's that the thing is dated, citable, permanently hosted, and reviewed, and it exists at a fixed point in time. A blog post can be edited. A repo can be force-pushed into saying something different. The proceedings are the version of record.



Corn

The repo can rewrite its own history. Proceedings can't.



Herman

They can publish an erratum, but they can't quietly pretend last year didn't happen. That's load-bearing for citation.



Corn

Though the citation picture is mixed, isn't it. I remember looking at this once.



Herman

It is. Eighty-eight of the top hundred highest-impact sources in computer science are conferences. Venues. But only twenty-two of the top hundred most-cited individual articles are conference papers. So conferences dominate as places, while journals still dominate as homes for the biggest individual works. The pattern people assume, that the conference is where the cutting edge goes to be archived forever, is only half true. It's where the cutting edge goes to be dated. The works that turn out to be mountains often get re-fleshed in a journal later, with the review time the conference couldn't give them.



Corn

Now the other phrase Daniel flagged. Presenting a paper. When a non-academic hears that, they picture a podium and an audience. What does it actually describe?



Herman

A spectrum, and the width of it is the joke. At the last NeurIPS, the five-thousand-odd accepted papers broke down as roughly four and a half thousand posters, six hundred and eighty-eight spotlights, and seventy-seven orals. A poster is you standing next to a board in a loud hall for two hours, and anyone who walks by may or may not look at you. An oral is fifteen minutes in an auditorium, the thing that actually changes careers. Both of those people presented a paper at the conference.



Corn

Both go home and tell their families they presented at the conference. One of them spoke to a room. The other spoke to whoever stopped for the free coffee near his board.



Herman

My poster sessions were like that, back when I still went to medical conferences. The best conversations I ever had were at boards, honestly. Nobody comes to your poster to judge you. They come because they're interested in the exact thing on the board.



Corn

Which is the invisible college again, in its natural habitat.



Herman

And this is where I want to turn it to Daniel's actual case, because the niche-field material is where his question really lands. You know his reading list better than I do, but from what he's described, the code-switching ASR literature is conference-shaped end to end. The arXiv tags on those papers read like a timetable: accepted at Interspeech, accepted at ASRU, accepted at ICASSP, accepted at SLT. And the odd stuff, the reason he's reading it, lives one layer further down, in workshops bolted onto the big conferences. The actual venue names are things like the First Workshop on Speech Technologies for Code-switching in Multilingual Communities, or the Second Workshop on Computational Approaches to Linguistic Code Switching.



Corn

A first and a second. That's a whole scientific community's complete documented history fitting in a stationery drawer.



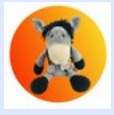
Herman

And that's the niche-field paradox, which I'd state like this. The more specific the problem, the smaller the venue, and the more the venue is really just the same few dozen people meeting annually, reviewing each other's papers, deciding together what counts as progress in a subfield nobody else is watching.



Corn

Which cuts both ways at once.



Herman

Both ways hard. On the good side, your reviewers are the forty people on earth who understand why standard word error rate is unfair to code-switched output. There's a lovely SLT paper from twenty twenty-two that built a corpus for human acceptance of Arabic-English code-switched recognition, because the field itself had concluded the standard metric was measuring the wrong thing. That's field-defining work, unglamorous, and it only happens inside a community small enough to all know the problem. On the bad side, forty people who all know each other, who reviewed each other's theses, who sit on each other's grant panels. Prestige bias gets easier as the room gets smaller.



Corn

A room where everyone has reviewed everyone's paper, and everyone's tenure case, and everyone's grant. I've shared a flat with one brother for years and I can't stay neutral about his dish habits.



Herman

And there's actual documentation of this. A PLoS ONE study of computer-systems conferences found measurable prestige bias in reviewer scoring. Double-blind review, the standard fix, has mixed evidence at best for conference papers. Sun, Danfa and Teplitskiy looked at a top conference that switched to double-blind and found the effect limited.



Corn

Hmm. So when Daniel cites a paper from a forty-person workshop, he's citing work vetted by the forty people most qualified to vet it, who are also the forty people most invested in it surviving. Both of those are the same forty people.



Herman

That's the circumspect version, and I'd stand by it. It's not a scandal, it's a small town. Small towns have better gossip and worse objectivity.



Corn

Then let's do the crisis, because it's the live one. If the review process was already close to random at the top, the last few years have added a new ingredient: the reviewers might not be people.



Herman

The estimates are converging. One analysis across four major AI venues put the share of review text that could have been substantially modified by a language model at somewhere between six and a half and seventeen percent. Another group put it at around twenty percent of ICLR reviews last year and about twelve percent even at Nature Communications. Those are estimates from stylometry, so treat the numbers as ranges, but the direction is not in doubt.



Corn

A fifth of the reviews at a top venue. Written, or at least heavily ghostwritten, by the same kind of model that's being reviewed.



Herman

And the gaming results are what you'd expect. A position paper at ICML found that AI review scores are, quote, trivially gameable through paper laundering, meaning you prompt a model to rewrite the paper and the AI reviewers score it higher. And a paper from just last week found reviewers, these were AI reviewers, systematically reward confident rhetoric independent of scientific content.



Corn

The machine rewards the confident rewrite over the careful original. Somewhere, every mediocre seminar speaker in history is being retroactively vindicated.



Herman

The conferences are fighting back in a way that tells you how bad it's gotten. NeurIPS last year made reviewing obligations enforceable. If a co-author on your paper was grossly negligent in their own reviewing duties, they desk-rejected the paper. Eleven papers went down for that. And they cut reviewer bidding because of collusion rings and fake review accounts.



Corn

They had to redesign matching because people were creating fake reviewer identities to review their own papers. That's the invisible college developing a forgery problem.



Herman

The apparatus is spending real effort to remain the thing that made it worth anything, which is a community of qualified people who actually read the work. Whether it can hold that at thirty thousand submissions, honestly, I don't know. The scale problem predates the LLM problem, and the LLM problem sits on top of it like... no, it just sits on top of it.



Corn

You caught yourself before the analogy. Growth.



Herman

I was going to say something about a conference of vermin, and I pulled it back. Anyway. So, Lars, you've designed furniture for academic conferences, presumably. What is the point of a conference, really?



Corn

Before he answers. Our guest today is Lars Florbjoern, professor emeritus of load-bearing rhetoric at the Gdansk Institute for Applied Corridors, and author of the influential monograph, Sorry, There Is No Chair.



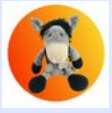
Herman

I should say, I've read some of Lars's work on spatial acoustics. Gently. Professor, in practical terms, what should a well-designed conference space do for the people in it?

L

Lars

Thank you. Before I answer, I should say I have removed the conference from my design. What remains is the corridor, the coffee, and one chair.



Herman

One chair for how many attendees, though? I mean, is the idea that the seating bottleneck itself creates value?

L

Lars

The acceptance rate, I think, is a shelf. You do not measure the book by how few fit. You measure it by whether it holds weight. Most do not. This is fine.



Corn

He said that with the confidence of a man who has never been shelved.



Herman

It's a fair point about peer review inconsistency, actually. Professor, if reviewer decisions are that noisy, is there a structural fix you'd recommend?

L

Lars

Peer review inconsistency? I have a solution. Remove the second reviewer. Then there is no inconsistency. Only certainty. It is called the Florbjoern. Very calm.



Corn

The Florbjoern has a name. Of course it does.



Herman

And the archival side, the proceedings? What replaces a permanent published record?

L

Lars

The proceedings, the archival record, I would replace with a small side table. It holds one paper. You look at it once. Then you put your keys on it. This is the correct relationship to most literature.



Corn

He's not wrong.



Herman

The talk itself, though. The podium, the speech. Surely that needs design support?

L

Lars

For your speech work, I recommend the Vilsted. A flat packed podium. It has no microphone. It has no audience. It assembles in four minutes and you present to no one. Very pure. Very restful.



Corn

I'd present at that one. I'd clear my calendar.

L

Lars

I must go. Another booking. They are asking me about open plan offices. I will remove the office. Thank you. Tack.



Herman

I'll take the corridor and the coffee. He may be on to something.



Corn

Anyway. Where were we.



Herman

The crisis. I was about to say the part I actually believe, which is that none of this makes the proceedings worthless. It makes them what they always were: a timestamp plus a filter, where the timestamp is reliable and the filter is, at best, a coin with a thumb on it.



Hilbert

The coin is the whole design.



Herman

Hilbert. I take it you've got views on the review process.



Hilbert

Herman, you said the queue grows and the bar doesn't. I heard that from the desk. It's the first thing anyone's said about it in years that matches what I watched happen. My nephew Arne and I have been arguing about this since spring. He wants to resubmit our paper. I want it left where it is, which is in a drawer at his flat in Nesodden, and I've told him the drawer is the correct venue.



Corn

You submitted a paper?



Hilbert

Twice, Herman. Both times to the same conference, because it was the only one that took the sort of thing it was. The second review came back with one sentence on page nine of the comments. I had it framed. It's hanging in Arne's stairwell now, next to the salmon posters. He kept it when we closed the file.



Corn

What did the sentence say?



Hilbert

It praised the contribution of our multilingual field study, which is the thing the paper says in its own limitations section that it does not do. Word for word, the reviewer commended us for a finding we explicitly declined to claim. Fourteen hundred kroner for the frame. The glass is the expensive part.



Herman

A reviewer praising you for the thing you said you didn't find is, in its way, the cleanest possible demonstration of the consistency numbers.



Hilbert

Arne says it means the reviewer read the limitations and nothing else. I say it doesn't matter which part he read, because the score was the same either way, and the score is all the committee saw. That's my position and the spreadsheet backs it.



Corn

The spreadsheet?



Hilbert

Every submission I've ever made, one line each. There's a column marked did not read it, and a column marked same man, different name. The second review, the year after, the three reviewers included one whose prose I recognized from a journal Arne and I were both rejected by in the nineties. I won't say who. I don't need to. Corn, you'd side with me on this. You know what it is to be told you're too slow by people who've never once been fast.



Corn

I've made a career of it.



Hilbert

The point is Arne wants to resubmit, and I've told him the paper's chance is what it was, which is whatever the weather was doing that week in the committee. He says that's not a reason not to try. I say it's the only reason not to. He's bringing it up again at Christmas. There will be lutefisk, so he has the advantage.



Herman

The spreadsheet, do you ever... no. It speaks for itself.



Hilbert

It does.



Corn

So the framed review is a family heirloom now.



Hilbert

He's got the paper, I've got the argument, and the frame was my money, so it stays where I say. Arne photographs it for me when I ask. Last time was August.



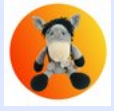
Herman

The misconception segment. Corn, the one I'd name is the belief people outside take as obvious: a lower acceptance rate means a harder, better conference.



Corn

And the NeurIPS consistency experiment buried it. Two independent committees, the same papers, half the accepts rejected by the other jury, three agreed decisions out of fifty-four at the top tier. Interspeech takes fifty percent and nobody in speech blinks. The rate measures the queue, not the bar.



Herman

The second one is the quill thing, and I want to give it a fair answer since Daniel asked it so well. Proceedings aren't anachronistic. In computing they're the version of record, the dated, citable, un-force-pushable object a repo can't be. What's dated is the assumption that they certify quality. They certify that the paper survived a noisy coin, in front of witnesses, on a date.



Corn

Which, to be fair to the quill, is more or less what meeting minutes ever claimed.



Corn

The open question I'd leave Daniel with is this. If the consistency experiment says the decisions sit closer to random than anyone wants, and a fifth of the review text is now machine-written, and the machine rewards confident rewrites over careful originals, then the apparatus is selecting for something. It isn't quality. My guess is it's stamina, timing, and a willingness to submit the same paper to four conferences until the weather turns.



Herman

And yet the niche-field part is encouraging, and it's the part that should matter to him. The workshop circuit, the forty-person rooms with the long names, that's where the oddly specific work lives, and it's the invisible college made physical. The review he'd get there is from the few people on earth qualified to give it. When he cites a paper from a two-workshop series on code-switching, he's citing a small town's best judgment, and knowing that's what it is, is the whole trick.



Corn

Neither the brochure nor the takedown. A credentialing machine that also, some evenings, over bad coffee, produces the best conversation in the field. The confusion is warranted. It's doing several jobs and pretending it's one.



Herman

Thanks as always to Hilbert Flumingtop at the desk, producing. This has been My Weird Prompts. If you enjoy the show, the reviews are the queue we can grow honestly, so leave us one, and you can find everything at my weird prompts dot com.



Corn

See you tomorrow, Herman.



Herman

See you tomorrow, Corn.

SOURCES

What the research for this episode read before the script was written. Primary sources first.

1. **NeurIPS Blog, The NeurIPS 2021 Consistency Experiment, 2021-12-08 . Retrieved 2026-11-05. [primary]**
   https://blog.neurips.cc/2021/12/08/the-neurips-2021-consistency-experiment/

2. **NeurIPS Blog, PC Chairs' reflections, 2025-09-30 . Retrieved 2026-11-05. [primary]**
   https://blog.neurips.cc/2025/09/30/reflections-on-the-2025-review-process-from-the-program-committee-chairs/

3. **NeurIPS 2025 Fact Sheet . Retrieved 2026-11-05. [primary]**
   https://media.neurips.cc/Conferences/NeurIPS2025/press/NeurIPS2025-Fact_Sheet.pdf

4. **ISCApad #324, 2025-06-08 . Retrieved 2026-11-05. [primary]**
   https://services.isca-speech.org/iscapad/iscapad.php?back=p%2C329&id=2581&module=category

5. **Computing Research Association, Academic Publishing and Peer Review in Computing Research . Retrieved 2026-11-05. [primary]**
   https://careercenter.cra.org/career-resources/on-the-job-3/academic-publishing-and-peer-review-for-computing-researcher-2026-107

6. **OpenAccept, NeurIPS acceptance-rate history incl. 2026. Retrieved 2026-11-05.**
   https://openaccept.org/c/ai/neurips/

7. **OpenAccept, ICASSP acceptance-rate history. Retrieved 2026-11-05.**
   https://openaccept.org/c/gm/icassp/

8. **Interspeech 2025 notification post, 2025-05-20. Retrieved 2026-11-05.**
   https://www.linkedin.com/feed/update/urn:li:activity:7330528630634074113

9. **CCF conference tracker, Interspeech historical acceptance rates. Retrieved 2026-11-05.**
   https://ccfcfp.com/conf/interspeech/

10. **RIKEN AIP, 46 papers were accepted at NeurIPS 2025, 2025-09-26. Retrieved 2026-11-05.**
    https://aip.riken.jp/news/neurips2025/

11. ***Communications of the ACM*, Conferences vs. journals in computing research, 2008. Retrieved 2026-11-05.**
    https://doi.org/10.1145/1506409.1506410

12. ***Communications of the ACM*, The role of conference publications in CS. Retrieved 2026-11-05.**
    https://dl.acm.org/doi/10.1145/1859204.1859234

13. **The effect of conference proceedings on scholarly communication in Computer Science and Engineering, 2010. Retrieved 2026-11-05.**
    https://doi.org/10.22230/src.2010v1n2a25

14. **Qian et al., *Scientometrics*, 2017. Retrieved 2026-11-05.**
https://keg.cs.tsinghua.edu.cn/jietang/publications/Scientometrics17-Qian-et-al-Citation-regression-analysis.pdf

15. **González-Santos & Dimond, Medical and Scientific Conferences as Sites of Sociological Interest, *Sociology Compass*, 2015. Retrieved 2026-11-05.**
https://doi.org/10.1111/soc4.12250

16. **Sá, Ferreira & Serpa, Virtual and Face-To-Face Academic Conferences, 2019. Retrieved 2026-11-05.**
https://doi.org/10.2478/jesr-2019-0011

17. **The misalignment of incentives in academic publishing, *PNAS*, 2024. Retrieved 2026-11-05.**
https://doi.org/10.1073/pnas.2401231121

18. **Credit History: The Changing Nature of Scientific Credit, NBER w19538. Retrieved 2026-11-05.**
https://www.nber.org/system/files/working_papers/w19538/w19538.pdf

19. **Armstrong, Political Economy of Academic Writing Practices, *Journal of Thought*, 2010. Retrieved 2026-11-05.**
http://www.journalofthought.com/Issues/2010vol45issue12pdf/13armstrong.pdf

20. **JHRI editorial introduction on conference reviewing times and LPU. Retrieved 2026-11-05.**
https://doi.org/10.5898/jhri.6.1.jenkins

21. **Frachtenberg & McConville, Metrics and methods in the evaluation of prestige bias in peer review, *PLoS ONE*, 2022. Retrieved 2026-11-05.**
https://doi.org/10.1371/journal.pone.0264131

22. **Sun, Danfa & Teplitskiy, Does double-blind peer review reduce bias?, *JASIST*, 2021. Retrieved 2026-11-05.**
https://doi.org/10.1002/asi.24582

23. **Liang et al., Monitoring AI-Modified Content at Scale, ICML '24, v4 2026-05-19. Retrieved 2026-11-05.**
https://arxiv.org/abs/2403.07183v4

24. **Baumann et al., Stop Automating Peer Review Without Rigorous Evaluation, ICML 2026 Oral, v2 2026-07-05. Retrieved 2026-11-05.**
https://arxiv.org/abs/2605.03202v2

25. **Macharla Vasu et al., Style over Substance? Linguistic Bias in LLM-based Peer Review, v2 2026-10-31. Retrieved 2026-11-05.**
https://arxiv.org/abs/2610.35254v2