The editor's signature: a proposal for AI-born journals

Family Name	Floridi
Given Name	Luciano
Division	Digital Ethics Center
Organisation	Yale University
Address	New Haven, CT, USA
Division	Department of Legal Studies
Organisation	University of Bologna
Address	Via Zamboni 27, 40126 Bologna, Italy
Email	luciano.floridi@yale.edu

Abstract

Since 2001, the online submission apparatus surrounding the world's scientific journals has absorbed an estimated one billion researcher-hours of cumulative labour. Half a million person-years. If a single researcher had been sitting at the same Editorial Manager submission screen, without sleep or interruption, since the height of the Last Glacial period, she would still be uploading supplementary files. The world has spent the equivalent of the entire direct workforce of the Human Genome Project, twice over, on filling in web forms. In this article, I ask whether it needed to and I argue that it did not, and that the technology to stop it has now arrived. The argument has a philosophical spine. Among the four (or five, depending on the literature) functions that hold the scholarly-communication system together, the journal's distinctive contribution is one: *certification*, performed by the editorial signature on a credible claim, namely, the act in which one or more accountable editors, embedded in a research community, backed by an institution, accept responsibility for what is published under conditions of evidential uncertainty. Everything surrounding that act is infrastructure, and AI is the first technology since the emergence of the modern journal apparatus that can absorb the infrastructure without disturbing the signature. Therefore, in this article I propose a new class of *AI-born* journals, defined by several necessary architectural features and governance controls. *Philosophy & Technology* could be a controlled experiment.

Keywords. scholarly publishing; artificial intelligence; AI-born journals; editorial signature; peer review; philosophy of information.

1. A quarter of a century is long enough

On 6 March 1665, Henry Oldenburg, a German diplomat then living in London and serving as one of the Royal Society's first two secretaries, posted from his lodgings the inaugural issue of the *Philosophical Transactions: Giving Some Accompt of the Present Undertakings, Studies, and Labours of the Ingenious in Many Considerable Parts of the World*.[1] The volume ran to sixteen pages, contained ten short pieces—including Robert Boyle's report on a monstrous calf and an account of a 'spot in one of the belts of Jupiter' communicated by Hooke—and sold for one shilling. There was no submission system. There was no peer review apparatus in the modern sense. There were no metadata fields, no ORCID identifiers, no conflict-of-interest declarations, no data-availability statements, no responses to reviewer comments, no formatting requirements, no preflight checks, no consent screens. There was Oldenburg, his correspondents, his judgement about what was worth circulating, and the printer.

Three hundred and sixty-one years later, an author submitting a manuscript to a representative mainstream scientific journal will spend between fifteen and eighteen hours of formatting and submission labour to see her paper through to publication.[2] She will reformat the body of her manuscript into the journal's house layout. She will reformat her references into the journal's preferred citation style. She will reformat her figures into the required resolution and file types. She will enter her co-authors' names, ORCIDs, and institutional affiliations into a web form that, for reasons no one alive can now reconstruct, requires each field to be entered twice. She will draft a cover letter. She will suggest reviewers and declare conflicts of interest. She will complete funding, ethics, and data-availability statements. She will upload her files in the order the system requires, and respond to the automated rejection of files that fail preflight.

[1]Oldenburg's correspondence has been edited and published in thirteen volumes by A. R. Hall and M. B. Hall (Madison: University of Wisconsin Press, 1965–1986). The first issue of the Philosophical Transactions can be consulted in facsimile in the Royal Society's digital archive at https://royalsocietypublishing.org/journal/rstl.

[2]The fifteen-to-eighteen-hour figure is my own estimate, derived by combining the LeBlanc et al. (2019) empirical median of fourteen hours of formatting labour per published manuscript through acceptance with an additional three-to-four-hour interface-labour uplift for cover letters, suggested-reviewer and conflict-of-interest declarations, funding and ethics statements, and residual interface friction. LeBlanc does not measure the interface uplift directly; the three-to-four-hour figure is a best-estimate triangulation. Section 2 sets out the derivation in detail.

She will consent to a sequence of screens whose legal weight no one has explained to her. Then she will click submit. If her paper is rejected, as most submissions to respected journals are,[3] she will repeat the performance at the next venue, with a new house style, a new form, a new set of preflight constraints. Some journals (the Royal Society portfolio among them) have begun to offer format-free initial submission, which reduces the first-round burden; the apparatus nevertheless reasserts itself at revision and typically reappears in full at the next venue tried.

No one intended this. The apparatus grew the way apparatuses grow: one reasonable addition at a time. The formatting requirements began as print-technology constraints in the late nineteenth century and survived into the digital era as proxies for craft. The metadata fields were added one by one to solve real problems (plagiarism detection, institutional accountability, interoperability with bibliographic databases), each addressing a specific need. The conflict-of-interest declarations responded to a genuine crisis of trust. The data-availability statements track a commitment to reproducibility. Each accretion, at the moment of its introduction, answered a real need. This is how any long-lived institution acquires its procedures; it is how a recipe handed down through generations acquires its ingredients until the dish no longer resembles the original. Nevertheless, the cumulative result is not what the recipe was for.

Editorial Manager launched in 2001.[4] ScholarOne, its older cousin, followed on from earlier Manuscript Central software and established itself as a journal infrastructure around the same time.[5] In 2026 Editorial Manager is twenty-five years old; ScholarOne, which marked its own twenty-fifth anniversary in 2025, is about twenty-six. A quarter of a century is long enough for the original design choices to

[3]Rejection rates at comparably-ranked scientific journals typically run in the range of 50 to 75 per cent, with Nature and the Nature-branded journals running substantially higher. Journal-level rejection rates are not systematically published across the field, but the aggregate picture is well documented in the scholarly-publishing literature and in the annual reports of several major publishers.

[4]Editorial Manager was developed by Aries Systems Corporation and launched in 2001; it was acquired by Elsevier in 2018. The Aries Systems company history and product timeline are documented at ariessys.com/about/history.

[5]ScholarOne Manuscripts was developed by ScholarOne Inc. as a successor to its earlier Manuscript Central product, first deployed in the late 1990s. The product was acquired by Thomson Reuters in 2006, then passed to Clarivate Analytics with the 2016 spin-off, and was acquired by Silverchair from Clarivate in late 2024.

have drifted out of living memory, for the procedures they encode to have been mistaken for the function they were meant to serve, and for the technology that would make a different architecture possible to arrive. All three have happened. The question this article addresses is what to do about it.

The argument in what follows has three parts. First, I want to show how much human labour the current apparatus now consumes. The number turns out to be large enough that the case for doing something does not depend on any claim that the apparatus is badly designed (it was designed well, twenty-five years ago) or that publishers are adversaries (they are not; they are partners who also bear the costs the current pattern imposes). Second, I want to sketch what a journal architecture built today from first principles would look like. I call this an *AI-born* journal: not a journal with AI bolted on, but one whose core workflow was conceived with what AI agents can reliably do in mind.[6] Finally, I want to argue for why this proposal should be attractive to publishers, not despite their commercial interests but because of them, and why the transition it describes is already underway, with or without their participation. I invite Springer-Nature, as the publisher of the journal you are reading, to lead it rather than respond to it.

2. Three clocks

Two products dominate the contemporary submission infrastructure. ScholarOne Manuscripts, acquired by Silverchair from Clarivate in late 2024, is reported by Silverchair to be used by more than 9,000 journal sites worldwide and to process around 3 million manuscripts per year, accounting for over 20 per cent of the world's scholarly journals.[7] Editorial Manager, built by Aries Systems in 2001 and owned since 2018 by Elsevier, is the other behemoth: it runs most Springer-Nature and Elsevier

---

[6] I introduce the term 'AI-born' as a structural category, not a marketing one. The distinction I intend to mark is between a journal that retrofits AI capabilities onto an inherited workflow (which I call AI-assisted) and one whose core architecture is conceived around what AI agents can reliably do (which I call AI-born). The latter, I argue, avoids the failure mode in which compressed per-task labour is offset by an expanded number of tasks. The wider design space for peer-review innovation is mapped in Tennant, J. P. et al. (2017), 'A multi-disciplinary perspective on emergent and future innovations in peer review', F1000Research 6: 1151; the AI-born proposal is one position within it.

[7] ScholarOne client figures from the Silverchair acquisition announcement, October 2024; product scale figures from ScholarOne's public documentation prior to the acquisition.

titles (including the journal in your hands), much of medical publishing, and many learned-society journals.[8] Below those two, eJournalPress (acquired by Wiley in 2021, used by PNAS among others), Open Journal Systems (the free PKP platform that runs more than 44,000 journals and dominates Latin American and small open-access publishing), and a handful of proprietary systems account for most of the remainder. Between them, ScholarOne and Editorial Manager handle a very large share of submissions to indexed journals. They are the apparatus.

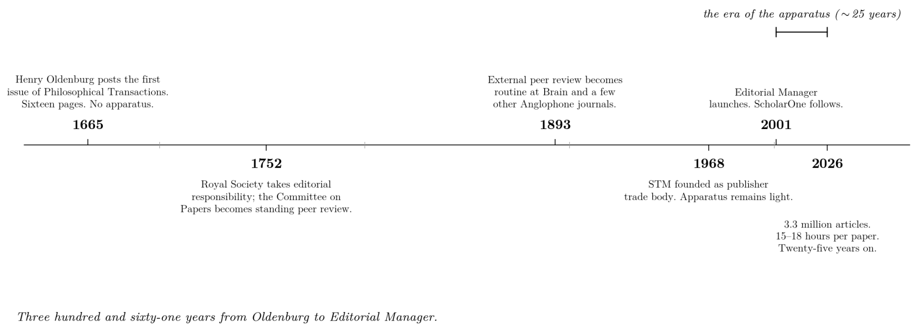

Figure 1. Three hundred and sixty-one years from Oldenburg to Editorial Manager. The submission apparatus, in its current form, is twenty-five years old; the journal it serves is three hundred and sixty-one.

The apparatus runs on *three distinct clocks*, and the rest of the article's argument depends on quantifying each of them honestly. The first is the labour the author performs to feed the apparatus: the hours she spends on formatting, on filling in web forms, on assembling submission packages. The second is the labour the editorial system absorbs to move a paper through it once submitted: the hours that reviewers, editors, editorial-office staff, copy editors, typesetters, and production teams collectively contribute to each paper that passes the gate. The third is the elapsed time the paper itself waits between submission and publication. The first two are measured in hours of human work; the third is measured in days of suspended scientific dissemination, which I will convert to a person-year-equivalent figure for comparability. The three are

[8]Editorial Manager market share is harder to estimate precisely, because Aries Systems does not publish a complete client list, but Springer-Nature, Elsevier, Wolters Kluwer, and the major medical-society publishers all use it as standard infrastructure.

dimensionally distinct, structurally connected, and, in their own ways, are now growing worse.

Before turning to the numbers, I should be explicit about what kind of numbers they are, because the case I am making depends on the reader trusting the construction. The three subsections that follow rest on different evidential foundations. Some figures are directly measured in published studies: LeBlanc et al. (2019) for author hours, Aczel et al. (2021) for reviewer hours, Andersen et al. (2021) and the Petrou series (2022–2025) for queue time. Some are bottom-up estimates I construct from this article's own architectural reasoning, which I show explicitly. Some are order-of-magnitude inferences in domains where no good source exists, and which I label as such. The body of each subsection identifies the type of figure each number represents. Where a source gives a range, I use a defensible mid-point—not the upper bound, which would inflate the case, and not the lower bound, which would understate it—and I show both extremes in the footnote so that any reader who prefers a different position in the range can substitute it. There are also moves I deliberately do not make, and which I name where they would otherwise tempt the writer: I do not extend the editorial clock to a cumulative-since-2001 figure, because the longitudinal data does not support it; I do not sum the three clocks into a single mega-number, because they measure dimensionally distinct quantities; and where the choice between a higher and a lower defensible figure is genuinely close, I take the lower. The intent is that the reader should be able to reconstruct each number from the page alone and find the conclusion robust to reasonable disagreement at every step.

2.1 The first clock: the author's time

How much labour does the apparatus extract per published paper from the author? The empirical anchor is LeBlanc et al. (2019), who surveyed 372 researchers across 41 countries and found a median of 14 hours of formatting labour per published manuscript from acceptance, with an interquartile range of 5 to 23 hours and a median of 2 submission attempts before acceptance.[9] The LeBlanc figure is for formatting

[9]LeBlanc, A. G., Barnes, J. D., Saunders, T. J., Tremblay, M. S., & Chaput, J.-P. (2019), 'Scientific sinkhole: The pernicious price of formatting', PLOS ONE 14 (9): e0223116. The study sampled 372 researchers from 41 countries via electronic self-report (October 2018 to January 2019), with snowball

work (body text, figures, tables, supplementary files, references) and excludes substantive writing, editing, and analysis. It does not capture the full time spent inside the submission interface; that broader figure has not been systematically studied. A narrower extrapolation by Jiang and colleagues, focused on biomedical journals, estimates 23.8 million researcher-hours per year worldwide on reformatting alone.[10] I take the LeBlanc fourteen-hour figure as the central per-manuscript estimate, with an explicit acknowledgement of its methodological limits and of the wide interquartile range. To it I add an estimated three to four hours of interface labour per accepted manuscript that LeBlanc does not directly measure (the cover letter, the suggested-reviewer and conflict-of-interest declarations, the funding and ethics statements, and the residual friction of consent screens, dropdowns, and navigational dead-ends), taking the total to between fifteen and eighteen hours of researcher labour per published paper.

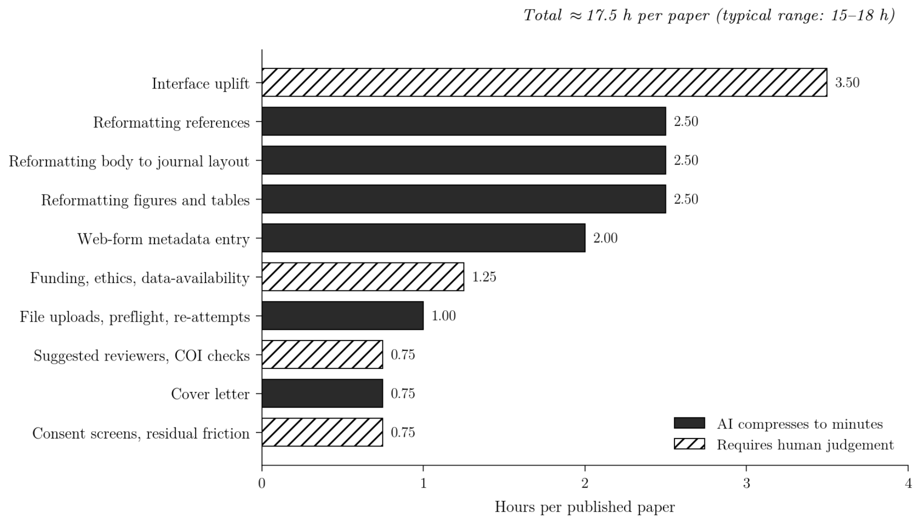

Where the fifteen-to-eighteen hours go. Decomposition follows LeBlanc et al. (2019) plus a 3.5-hour interface uplift.

sampling that produced a 60-per-cent Canadian skew. The reported 14-hour median is for formatting labour per published manuscript through acceptance, not per submission attempt. Interquartile range: 5 to 23 hours.

[10] Jiang, Y., Lerrigo, R., Ullah, A., Alagappan, M., Asch, S. M., Goodman, S. N., & Sinha, S. R. (2019), 'The high resource impact of reformatting requirements for scientific papers', PLOS ONE 14 (10): e0223976.

Figure 2. Where the fifteen-to-eighteen hours go. Solid black bars indicate components that current AI tools compress to minutes; diagonally hatched bars indicate components that require human judgement. Decomposition follows LeBlanc et al. (2019) plus a 3.5-hour interface uplift.

Figure 2 decomposes the per-manuscript burden. The solid black bars (reformatting body text, references, and figures, plus web-form metadata entry, file uploads, and cover letter drafting) total roughly twelve hours and are each, with current commercial tools, a problem AI compresses to minutes. The diagonally hatched bars (suggested reviewers and conflict-of-interest checks, funding and ethics statements, residual friction) total roughly four to five hours and require human judgement that AI lacks. Two-thirds of the per-paper burden, on this accounting, is technically compressible. One-third is not. The compressible two-thirds is what section 5 will argue AI can absorb. The irreducible one-third marks the human-attestation layer on the author side that the architecture must preserve rather than automate; the journal's own irreducible act lies elsewhere, in the editorial synthesis defended in §7.

To turn the per-paper figure into a global annual aggregate, multiply by the world's annual published output. Indexed-journal article counts vary. Clarivate's Web of Science indexed roughly 3.3 million scientific articles in 2024, of which about 2.5 million were classified as research studies; broader counts that include conference papers, reviews, and editorials reach 5 million or more.[11] Using the Clarivate figure of 3.3 million as the conservative denominator, and the fifteen-to-eighteen-hour per-paper estimate, the global annual labour absorbed by formatting and the submission interface runs to between 50 and 60 million researcher-hours. To this should be added the labour absorbed by submissions that never publish anywhere (papers rejected at

[11] Article-output figures also vary. Clarivate's Web of Science indexed approximately 3.3 million scientific articles in 2024, of which about 2.5 million were classified as research studies; broader counts that include conference papers, reviews, and editorials reach 5 million or more. I use the Clarivate figure of 3.3 million as the conservative denominator. The dark-submission uplift accounts for papers that are formatted, submitted, rejected at every venue they reach, and never publish anywhere; I estimate this at 30 to 50 per cent of the published flow, but the figure is an informed assumption rather than a measured total: the dark-submission population is not systematically tracked. For this reason, the cumulative and annual labour figures in this section that incorporate the uplift (70 million hours per year; 1 to 1.3 billion cumulative hours) should be read as scenario estimates. The LeBlanc-only lower bound, roughly 700 million hours or 350,000 person-years without any uplift, is the defensible empirical floor, and the argument of the article stands on that floor alone.

every venue they reach), which plausibly raises the figure by 30 to 50 per cent, to 65 to 90 million hours per year. I use 70 million hours per year as the central working estimate, with 50 million as the LeBlanc-only floor and 90 million as the dark-submission ceiling.

Cumulatively, since Editorial Manager launched in 2001, the apparatus has been with us for twenty-five years. Indexed-journal article output rose from approximately 1.1 million in 2001 to 3.3 million in 2024, a compound annual growth rate of around five per cent. Integrating the curve gives roughly 50 million accepted papers across the period. Multiplied by the central per-paper figure, with the dark-submission uplift, the cumulative aggregate falls in the range of 1 to 1.3 billion researcher-hours. At a 2,000-hour working year, this is the equivalent of 500,000 to 650,000 person-years of dedicated researcher labour. The lower bound, using only the LeBlanc fourteen-hour median with no uplifts, gives roughly 700 million hours, or 350,000 person-years. The honest cumulative range is therefore 700 million to 1.3 billion hours, equivalent to between 350,000 and 650,000 person-years. Around half a million person-years is the central figure for the first clock alone, and I use it as the author-side working estimate for the rest of this article.

2.2 The second clock: the editorial system

The first clock counts only authors. The apparatus extracts labour from the other side of the desk as well, and on a larger scale than the author-side figure suggests. The second clock measures the cumulative human work the editorial system absorbs to move papers through it once they have been submitted: the time that reviewers, editors, editorial-office staff, copy editors, typesetters, and production teams contribute to each paper. I will quantify it in the components below, with the evidential status of each named explicitly.

The firmest leg of the calculation is the reviewer side. Aczel, Szaszi, and Holcombe (2021) constructed a global estimate of reviewer time donated to peer review by combining mean review counts per paper, mean hours per review, and global publication volumes. They estimate approximately 130 million hours of reviewer labour donated globally in 2020: 'a billion-dollar donation', as they title their paper.[12] The estimate has methodological caveats: it triangulates from survey data and publisher reports rather than measuring directly, and Aczel and colleagues themselves note that their figure is likely an underestimate because it covers only a portion of the world's journals. It is nevertheless the strongest single quantification of the reviewer side that exists, and I use the published figure of 130 million reviewer-hours per year as the central estimate. For the floor I use 100 million, which is the figure Aczel and colleagues report in their abstract; for a plausible ceiling, 160 million as a modest upward adjustment for journal coverage gaps. The bracketing figures are mine; the central 130 million is theirs.

The editor and editorial-office side is harder. No global figure has been constructed for it, and the publishers do not release operational labour figures by paper. I therefore construct a bottom-up estimate from this article's own architectural reasoning, and I want to be transparent about how the construction runs. In section 6 I posit that an editor processing a paper through the traditional pipeline spends approximately two hours on reviewer-matching and the invitation cascade, plus three hours on synthesising the reviews and writing the decision letter, which is roughly five hours of substantive editor time per paper that goes through review. Multiplied by the 3.3 million indexed-paper denominator, this gives approximately 16 million editor-hours per year on the published flow alone. Adding the editor time spent on rejected submissions (papers that are desk-rejected or rejected after review at one or more venues before publishing elsewhere) brings the figure to roughly 20 to 25 million editor-hours per year.[13] To this must be added the labour of editorial-office staff

[12] Aczel, B., Szaszi, B., & Holcombe, A. O. (2021), 'A billion-dollar donation: Estimating the cost of researchers' time spent on peer review', Research Integrity and Peer Review 6: 14. The authors estimate approximately 130 million hours donated to peer review in 2020, with US-based reviewers' time alone valued at over $1.5 billion in foregone professional time. The methodology builds on Kovanis, M., Porcher, R., Ravaud, P., & Trinquart, L. (2016), 'The global burden of journal peer review in the biomedical literature: strong imbalance in the collective enterprise', PLOS ONE 11(11): e0166387, the precursor estimate for the biomedical literature alone.

[13] The bottom-up editor-time estimate runs as follows. In §6 I posit that an editor in the traditional pipeline spends approximately two hours per paper on reviewer-matching and the invitation cascade, plus three hours on synthesising reviews and writing the decision letter: roughly five hours of substantive editor time per paper that goes through review. Multiplied by 3.3 million indexed-paper denominator, this gives approximately 16 million editor-hours per year on the published flow alone. Adding editor time spent on rejected submissions (papers desk-rejected or rejected after review at one or more venues

(managing editors, editorial assistants, integrity-check personnel) who are not the academic editor but support the workflow. I have no clean source for this figure. As an order-of-magnitude inference, if the global editorial-office headcount is in the range of 20,000 to 30,000 full-time equivalents at 2,000 working hours per year, this contributes a further 40 to 60 million hours per year to the second clock.

The production side—copy editors, typesetters, proofreaders, production project managers—is harder still. I have no global labour figure for it, and the publishers I have asked do not publish one. The order-of-magnitude inference is that production labour per paper is in the range of five to fifteen hours, including copy editing, typesetting, proof handling, and integrity verification. At 3.3 million papers per year, this amounts to roughly 15 to 50 million production-staff hours per year. I take the lower end of this range—20 million hours per year—as the cautious central figure, since the higher end is more inferential than I am comfortable defending here.

The cautious sum of the second clock is therefore approximately 130 million reviewer-hours, plus 25 million editor-hours, plus 50 million editorial-office-staff hours, plus 20 million production-staff hours, for a total of approximately 225 million hours per year. The conservative floor (Aczel's reviewer figure alone, treating the editor and production figures as too inferential to bank on) is 130 million hours per year. The ceiling, using the upper bounds of each inferential extension, is approximately 300 million hours per year. I use 225 million hours per year as the central working estimate, with 130 million as the floor and 300 million as the ceiling, and I will use the central figure for the rest of this article while reminding the reader at the §3 pay-off that the conclusion holds even on the floor alone.

Two observations follow. First, the second clock is materially larger than the first. On the central estimates, 225 million editorial-system-hours per year against 70 million author-hours per year: a ratio of roughly three to one. Even on the floors—130 million against 50 million—the ratio is approximately two and a half to one. The labour the apparatus absorbs from the editorial system is two to three times that absorbed from authors, which is itself a conservative reading of a body of evidence that may understate both. Second, the editorial-system labour is structurally different from the

before publishing elsewhere) brings the figure to roughly 20 to 25 million editor-hours per year. The 25 million figure is used as the central in body. The figure is bottom-up reasoning, not a measured total.

author-side labour: it is largely unpaid academic service for reviewers and editors, and a paid publisher-side cost for editorial-office and production staff. The compression argument I will develop in section 5, therefore, lands on different parties for different components: section 6 will return to this when discussing how the AI-born architecture affects each.

I want to explicitly flag the two things I am not doing in this subsection, because both would be tempting and would inflate the case. I am not extending the second-clock figure to a cumulative total since 2001. The author-side cumulative calculation works because LeBlanc's 14-hour figure is reasonably stable across the 2001–2024 window, and submission counts are tracked. The second clock has no equivalent longitudinal anchor: Aczel is one snapshot in 2020, the editor and production figures are inferential extensions, and the longitudinal evidence available (the Petrou series in the next subsection) is for elapsed time rather than total hours. Constructing a cumulative figure since 2001 for the second clock would require assumptions stronger than the data support. I am also not summing the first and second clocks into a single combined annual figure, even though the temptation to add 70 + 225 = roughly 300 million hours per year is real. The two are commensurable in units (both are hours of human labour), but their evidential status differs sharply, and combining them into a single number obscures the distinction the reader needs in order to evaluate the argument. I therefore report them separately throughout.

2.3 The third clock: the queue

The first two clocks measure human labour. The third clock measures something different: the elapsed time a paper spends waiting between submission and publication while the system works on it (or, more often, while it sits in someone's queue waiting to be worked on).[14] Where the labour clocks are measured in hours per paper, the

[14]The queue clock measures the elapsed wall-clock time a paper spends between submission and final publication. To make it commensurable with the labour clocks (which are measured in hours of human work per year), I convert it to a paper-year-equivalent: the number of paper-years held in suspension at any given moment across the global flow. The conversion is simply: mean submission-to-publication days × annual paper output / 365 days per year. On the central estimate of 250 days mean submission-to-publication and 3.3 million indexed papers per year: 250 × 3,300,000 / 365 ≈ 2.26 million paper-years. This is paper time, not human time; the labour clocks already account for the human contribution.

queue clock is measured in days, and where the labour figures aggregate across the global flow, the queue figure aggregates as person-year-equivalents of suspended scientific dissemination. The third clock is the one most visible to authors as frustration and most consequential for the velocity of science.

The current state of the queue clock is well documented across the recent literature. Andersen, Fonnes, and Rosenberg (2021) conducted a PROSPERO-registered systematic review of 69 studies on handling times in biomedical journals. Mean submission-to-publication times in the studies they reviewed ranged from 91 to 639 days, with median submission-to-publication times running from 70 to 558 days. Submission-to-acceptance ranged from 50 to 276 days; acceptance-to-publication from 11 to 362 days. Their conclusion: 'Editorial handling times of journals varied widely from a few months to almost two years, which delays the availability of new evidence.'[15] Chen et al. (2024), examining primary health-care journals, report an average submission-to-publication lag of 243 days, with peer-review and production phases of 178 and 66 days, respectively.[16] A 2025 cross-sectional study of 57 health policy journals reported median submission-to-publication times ranging from 35 to 353 days, with the peer-review phase reaching 314 days at the slow end. Björk and Solomon (2013) found cross-disciplinary averages of nine to eighteen months, with social sciences and humanities journals roughly double the chemistry and other STEM figures.[17] The honest summary is that, on the central estimate, the mean submission-to-publication time across indexed journals is in the range of 200 to 300 days, with

---

I am grateful to Jessica Morley for the observation that AI-born journals would also compress publication delays, which prompted the development of the third clock.

[15]Andersen, M. Z., Fonnes, S., & Rosenberg, J. (2021), 'Time from submission to publication varied widely for biomedical journals: a systematic review', Current Medical Research and Opinion 37 (6): 985–993. PROSPERO registration CRD42020196238. The systematic review covered 69 included studies on biomedical journal handling times. Mean submission-to-publication times in the studies they reviewed ranged from 91 to 639 days; medians from 70 to 558 days. Submission-to-acceptance: 50 to 276 days. Acceptance-to-publication: 11 to 362 days.

[16]Chen, T.-A., Lin, M.-H., Chen, Y.-C., & Chen, T.-J. (2024), 'The time from submission to publication in primary health care journals: a cross-sectional study', Publications 12 (2): 13. The 2025 cross-sectional study of health policy journals: see Phillips, K. A., & Horn, D. M. (2025), 'Review and publication times and reporting across journals on health policy', JAMA Network Open 8 (5): e2512545. doi:10.1001/jamanetworkopen.2025.12545.

[17]Björk, B.-C., & Solomon, D. (2013), 'The publishing delay in scholarly peer-reviewed journals', Journal of Informetrics 7 (4): 914–923. Their cross-disciplinary analysis found total submission-to-publication varying from approximately 9 months (chemistry, fastest) to 18 months (business and economics, slowest), with social sciences and humanities averaging 1.5 years.

high variance by field, and that several months of this are reviewer-and-editor wait time rather than substantive review time.

The longitudinal trend matters even more than the current state. Ellison (2002) documented that mean review times at top economics journals more than doubled between 1970 and 1999, from approximately 8.7 months to roughly 20 months on average, with individual journals tripling. Ellison attributed this primarily to the increasing number of iterative rounds of revision required, a structural cause rather than a circumstantial one.[18] The most useful recent longitudinal evidence comes from Christos Petrou's analyses of operational data from the major publishers, published in The Scholarly Kitchen between 2022 and 2025. In a 2022 study of more than 700,000 papers across 10,000 journals from the ten largest publishers, Petrou found that the mean total turnaround time fell from 199 days in 2011/12 to 163 days in 2019/20, a 36-day improvement over the decade. The reduction was driven primarily by production-side digitisation gains (XML workflows, online-first publication, automated typesetting), with peer-review compression a smaller component and concentrated in two publishers. In a follow-up 2024 piece, Petrou observed that the mainstream review process had remained roughly as slow as it was a decade earlier. In a 2025 piece based on a larger dataset of 8 million papers across 16 publishers, Petrou reported that submission-to-acceptance time in 2024 was the slowest since at least 2011, and nine days slower than in 2014. The reversal was near universal: 11 of 12 sampled publishers were slower in the second half of 2024 than in the first half of 2021.[19]

[18]Ellison, G. (2002), 'The slowdown of the economics publishing process', Journal of Political Economy 110 (5): 947–993. Ellison's Table 1 reports mean total review times at top economics journals across four decades. Examples (in months): Econometrica 8.8 in 1970, 14.0 in 1980, 22.9 in 1990, 26.3 in 1999. Review of Economic Studies: 10.9 in 1970 to 28.8 in 1999. Ellison attributes the slowdown primarily to increasing iterative rounds of revision required, a structural cause.

[19]Petrou, C. (2022), 'Publishing fast and slow: a review of publishing speed in the last decade', The Scholarly Kitchen, 8 November. Petrou, C. (2024), 'Publishing fast or slow: how speed varies for similar journals', The Scholarly Kitchen, 19 March. Petrou, C. (2025), 'How the growth of Chinese research is bringing Western publishing to breaking point', The Scholarly Kitchen, 8 July. The 2022 piece is based on more than 700,000 papers across more than 10,000 journals from the ten largest publishers, comparing 2011/12 to 2019/20. The 2025 piece is based on 8 million papers across 16 major publishers (excluding MDPI). Petrou's framing of the trend as 'sleepwalking toward an avoidable breaking point' is in the 2025 piece.

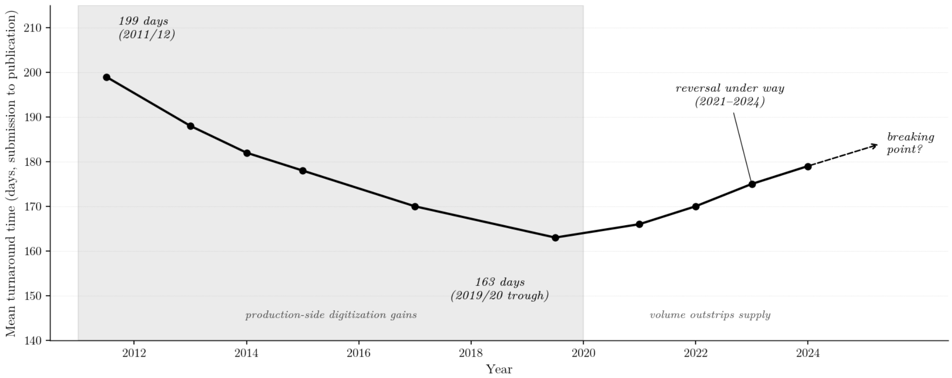

Figure 3. The U-curve of publishing turnaround time, 2011–2024. Production-side digitization drove a 36-day improvement in mean total turnaround time across the major publishers between 2011/12 and 2019/20 (199 to 163 days). The trend has since reversed: by 2024, submission-to-acceptance time across 16 major publishers was the slowest since at least 2011, with 11 of 12 sampled publishers slower in late 2024 than in early 2021. Data points reconstruct Petrou's published figures (2022, 2024, 2025); intermediate years are interpolated for visual continuity. The dashed projection beyond 2024 indicates trajectory, not measurement.

The diagnosis Petrou offers is structurally important to this article's argument. The decade of speedup, 2011–2019, was driven by one-time digitisation wins on the production side. Those gains have now been banked, and there is little remaining production speedup available without architectural change. The reversal since 2021 is driven by a volume-versus-supply mismatch: paper output is rising sharply while the editorial-and-reviewer pool is roughly static. Chinese paper output grew at 11.3 per cent annually from 2015 to 2024, from 464,000 to 1,216,000 papers per year, while the United States grew at 0.5 per cent annually, and most other developed-country output was largely flat. Researcher full-time-equivalent counts in the developed countries that supply most editors and reviewers grew at low single digits annually. Public-health journals saw 25 per cent growth in submissions from 2019 to 2020 and a further 21 per cent from 2020 to 2021. The 2016 Publons editor survey found that 75 per cent of editors identified 'finding reviewers and getting them to accept review invitations' as the most difficult part of their role. Petrou's 2025 forward projection: if these trajectories continue, the system will move toward what he calls a 'breaking point': slower review, more rejections without review, more papers diverted to faster but lower-quality venues, and progressive decoupling between the volume of submitted research and the volume of curated publication.

To make the queue clock comparable in scale to the first two, I convert the elapsed time to a person-year equivalent of suspended publication. On the central submission-to-publication estimate of 250 days across the global flow, with 3.3 million indexed papers per year, the queue holds approximately 250 days × 3.3 million papers ÷ 365 days, which is approximately 2.3 million paper-years of suspension at any given time. This is not 2.3 million person-years of human work—the labour clocks already account for that—but rather the volume of completed scientific findings sitting in editorial queues at any moment, awaiting public dissemination. Petrou's own related calculation: the nine-day slowdown observed since 2014, applied across the global flow, translates to approximately 80,000 paper-years of additional waiting time per year from the marginal increase alone. The floor estimate using the lower bound of Andersen's range (mean 70 days for the fastest journals) gives approximately 600,000 paper-years; the ceiling using the upper bound (mean 500 days for the slowest) gives approximately 4.5 million paper-years. I use 2.3 million paper-years as the central working estimate.[20]

Let me close this section with two observations. First, the third clock is structurally connected to the second. The queue is long because the editorial system is overwhelmed: editors take weeks to find willing reviewers because there are too few; integrity checks take days because copyeditors are overstretched; decision letters wait because senior editors are processing too many of them. Petrou's diagnosis of why the U-curve has reversed is directly a second-clock failure manifesting as a third-clock symptom. The two cannot be addressed independently. Second, all three clocks now point in the same direction: more author labour absorbed per paper, more editorial-system labour absorbed per paper, longer queues per paper, and the trend on each is rising. The labour cost is growing faster than the labour supply, the production gains

[20]The 80,000 paper-years figure is from Petrou (2025), calculated as the nine-day slowdown observed since 2014 multiplied across the global flow. The 2.3-million paper-year central figure is mine, calculated from the 250-day mean and the 3.3-million-paper denominator as set out in footnote 14. The floor of 600,000 paper-years uses 70 days as the mean (lower bound of the Andersen 2021 range); the ceiling of 4.5 million uses 500 days (upper bound). All three figures should be read as approximate orders of magnitude rather than precise totals.

have been exhausted, and no incremental adjustment to the existing architecture has reversed any of the three trends. This is the empirical situation that the proposal in §9 is responding to.

Three caveats close section 2. First, the labour costs fall disproportionately on the most productive researchers (those who submit most often, and who therefore reformat most often) and on the most senior editors and reviewers (those most often invited). The labour distribution across the three clocks is heavily skewed, and the per-individual burden is unevenly felt; the apparatus is regressive in the technical sense. Second, the figures throughout this section are global aggregates and conceal large per-discipline and per-region variation. A chemistry paper, in a mainstream journal, experiences something close to the floor estimate on each clock; a humanities or economics paper, in a slow journal, experiences something close to the ceiling. Third, the three clocks are growing worse independently rather than as a single coordinated phenomenon, but they share a common structural cause: the volume of papers entering the apparatus is rising faster than the apparatus can adapt, and the apparatus, in its current form, has no available adjustments left.

3. One billion hours, in human terms

The cumulative author-side figure from §2 is one billion hours, give or take. The annual editorial-system figure adds another 225 million hours per year on the central estimate, or 130 million on the conservative floor. The queue clock holds roughly 2.3 million paper-years of suspended dissemination at any moment. These are abstractions. Let me convert one of them.

I take the author-side cumulative figure for the conversion because it is the only of the three clocks for which the longitudinal data support a since-2001 calculation. A year contains 8,766 hours. One billion hours, divided by 8,766, is approximately 114,000 years. If we walk that span backwards from today (27 April 2026), we land somewhere around 112,000 BCE, deep in the Late Pleistocene. Anatomically modern *Homo sapiens* was by then already well established in Africa, with fossils dating to the Jebel Irhoud finds at roughly 300,000 years ago, and our species had begun the slow expansion that would eventually take it across Eurasia. Behavioural modernity (what we recognise as art, ritual, and complex symbolic culture) was still 40,000 to 60,000 years away. Agriculture would not appear for another 100,000 years. The first major city, Uruk, would not be founded for another 108,000 years. The Pyramid of Khufu sits roughly 109,000 years closer to us. The Roman Empire is one-fiftieth of the way back. Aristotle is, by comparison, a contemporary.

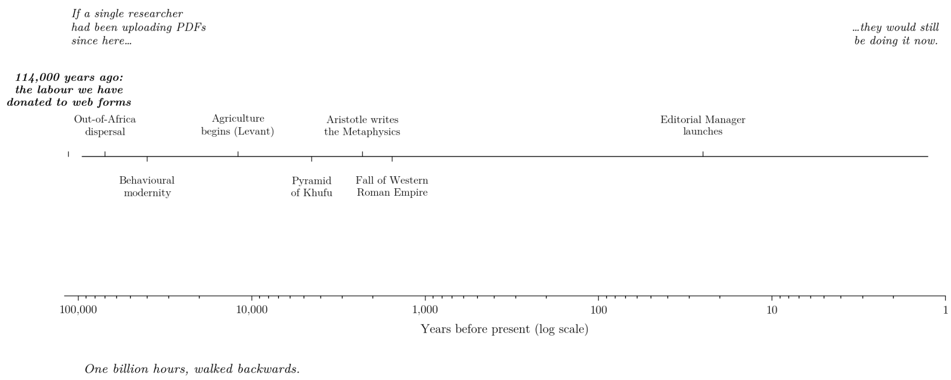

Figure 4. One billion hours, walked backwards. Divided by a working year of 8,766 hours, one billion hours equals roughly 114,000 years, which lands in the Late Pleistocene. The figure converts the author-side cumulative-since-2001 estimate from §2.1 only; the annual second-clock and third-clock figures from §2.2 and §2.3 are not included in this calculation, for the methodological reasons given in §2.

The image is rhetorical, but it is not exaggerated. If a single researcher had been sitting at the same Editorial Manager submission screen, without sleep or interruption, since the height of the Last Glacial period (when our species had reached the Levant but not yet Europe, and Neanderthals still walked the Eurasian steppe), she would still be uploading supplementary files. The Pharaohs would have come and gone in a rounding error. Aristotle would have written the *Metaphysics* in the time it takes to enter the co-author affiliations. The Library of Alexandria would have burned and been forgotten in the time it takes to fill in the funding statement. And, again: this is the author-side figure alone. The second clock, on the central editorial-system estimate, runs at roughly three times this rate every year on top, and the third clock holds another two million paper-years in suspension at any given moment.

For most readers, a more useful comparison will be against contemporary scientific projects whose labour costs are at least roughly known. Figure 5 sets the submission apparatus against four landmark efforts. The Apollo programme, on Charles Fishman's well-known accounting, involved 410,000 people across thirteen years of dedicated effort.[21] The Human Genome Project, on the Battelle economic-impact study, generated approximately 711,000 direct job-years between 1988 and 2010, with total employment impact, including indirect and induced effects on the genomics-enabled industry, reaching 3.8 million job-years.[22] The Higgs boson (from Brout, Englert, and Higgs's theoretical proposal in 1964 to the experimental discovery in 2012) required an order of magnitude less direct effort, with thousands of physicists working at CERN over four decades. LIGO, from Rainer Weiss's initial 1972 work to the 2015 detection of gravitational waves, required less still. None of these projects has a single canonical labour figure, and the comparisons in Figure 5 should be read as illustrative orders of magnitude rather than precise accounting.[23]

---

[21] Fishman, C. (2019), One Giant Leap: The Impossible Mission That Flew Us to the Moon (New York: Simon & Schuster). Fishman's figure of 410,000 people across the Apollo programme is the most widely cited but counts peak workforce rather than person-years. For an improved cost analysis, see Dreier, C. (2022), 'An improved cost analysis of the Apollo program', Space Policy 60: 101476.

[22] Tripp, S., & Grueber, M. (2011), Economic Impact of the Human Genome Project (Columbus, OH: Battelle Memorial Institute). The Battelle study reports 711,000 direct job-years for the Human Genome Project for 1988–2010, with a total employment impact of 3.8 million job-years across direct, indirect, and induced effects over the same period.

[23] Person-year figures for major scientific projects are not standardised and vary widely depending on whether one counts contractors, partial-time involvement, spillover employment, or only dedicated technical effort. The comparisons in Figure 5 should be read as illustrative orders of magnitude rather than precise labour accounting. For the Higgs discovery, see Aad, G., et al. (2015), Physical Review Letters 114 (19): 191803; no single canonical labour figure exists for the collaboration over its decade-long trajectory.

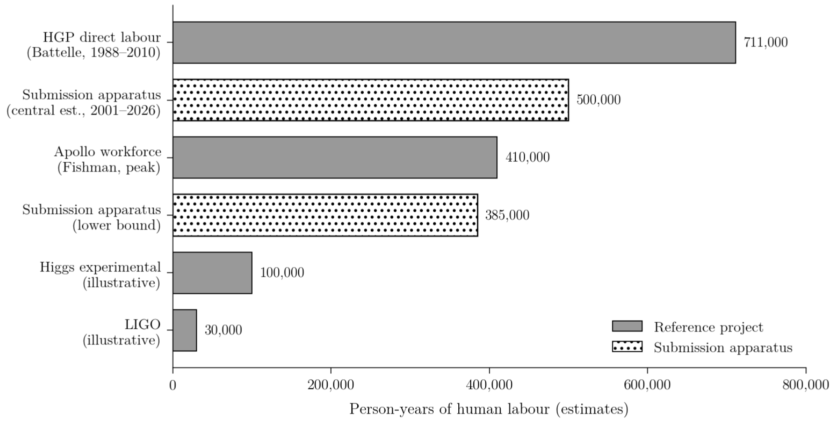

Submission apparatus against four landmark projects (illustrative orders of magnitude).

Figure 5. First clock against four landmark projects (illustrative orders of magnitude). Gray bars are reference projects; stippled bars are the submission apparatus first-clock cumulative-since-2001 figure (central and lower-bound estimates). Project figures use different definitions (workforce, direct job-years, dedicated effort); see footnote 23. Second-clock and third-clock figures are not represented here for reasons of dimensional commensurability.

Set the first clock against these benchmarks. On the central estimate of around 500,000 person-years, the author-side labour absorbed by Editorial Manager and ScholarOne since 2001 is comparable in magnitude to the direct labour of the Human Genome Project. It is larger than the Apollo workforce of 410,000 people at peak headcount, and it dwarfs the dedicated effort that produced the Higgs discovery and LIGO. Even on the lower-bound estimate of 350,000 person-years, the comparison holds: the submission apparatus has absorbed author-side cumulative labour of the same order of magnitude as the Human Genome Project. Apollo's labour bought us six manned lunar landings. HGP's labour bought us a complete reference sequence and the genomics industry. The submission apparatus's labour bought us the ability to process an ever-growing number of submissions to the journals that publish the science the other efforts produced: a real achievement, not a trivial one, and one that has served the community well at a certain scale. The question is whether it is still the best use of the community's time now that a different architecture is technically available, and whether the second and third clocks, which are growing worse rather than holding steady, will continue to be sustainable on the apparatus's current trajectory.

A reader with a sceptical disposition will object that the comparisons are not strictly comparable. The Apollo workforce concentrated its labour on a single mission with a defined output; the submission apparatus distributes its labour across millions of authors, editors, and reviewers, each spending hours per paper, in support of all 50 million articles published in the period. Concentrated mission-labour and distributed micro-labour are different kinds of costs, and the analogy partially deflates as a result. The objection is correct as far as it goes. But the residual question—what does the labour buy?—survives the deflation. Apollo's labour bought lunar landings. The submission apparatus's labour bought a function (the editor's signature on a credible claim) that, as I shall argue in section 6, can now be delivered with a fraction of the input across all three clocks. The asymmetry between what we spend and what we could spend, once the comparison is set out plainly, is what makes the proposal worth making.

4. The productivity paradox

The submission apparatus does not exist in a vacuum. It exists in a research economy that, on the best available evidence, has been losing productivity for decades. Bloom, Jones, Van Reenen, and Webb (2020), in what is probably the most cited recent paper on the economics of research, asked a deceptively simple question: how much output does a unit of research input produce, and is that ratio changing over time? Their answer, across a wide range of industries and at the aggregate level for the United States, is that research productivity has fallen by a factor of 41 since the 1930s. The famous Moore's Law doubling of transistor density now requires more than 18 times as many researchers as it did in the early 1970s.[24] Crop yields, drug-development pipelines, and aggregate total factor productivity all show the same pattern: stable or slowing output growth, sustained only by exponential increases in research effort.

---

[24]Bloom, N., Jones, C. I., Van Reenen, J., & Webb, M. (2020), 'Are ideas getting harder to find?', American Economic Review 110 (4): 1104–1144. The 41-fold productivity decline is the aggregate US figure; the Moore's Law 18-fold researcher multiplier is the domain-specific figure for semiconductor density.

The Bloom et al. paper has attracted criticism, and the criticisms have force. Output measures may not capture the value of new ideas; declining productivity may partly reflect compositional shifts in the research workforce; measuring 'researcher' is difficult across decades. Subsequent work has been influential and partly supportive, though the interpretation remains disputed: the existence of a slowdown is now widely accepted, but its magnitude, its causes, and whether it implies falling research productivity in the strong sense Bloom and colleagues argue are all live questions. The honest summary is that something real is happening; the size and mechanism are still being contested.

Bloom and his colleagues attribute the slowdown to a combination of two factors. First, the genuine difficulty of finding new ideas as the easy ones are used up (Ben Jones's 'burden of knowledge' hypothesis, in which the prerequisite training to reach the research frontier in any field has grown longer, leaving fewer productive years per researcher, a real effect, well documented).[25] Second, the rising complexity of research itself, which requires larger teams, more equipment, and more coordination. They do not name administrative friction as a cause. They do not name the submission apparatus, the formatting tax, the proliferation of compliance frameworks, the doubling of conference and grant-application overhead, or any of the other invisible taxes that have grown around the research enterprise in the same period. The omission is understandable (these things are hard to measure and harder still to attribute), but it is not innocent.

Consider the timing. The compounding of the productivity decline Bloom and colleagues document has accelerated since approximately 2000. In my accounting, the submission apparatus came into widespread use in 2001 and reached near-universal use by 2010. The formatting tax, the metadata tax, the conflict-of-interest-declaration tax, the data-availability tax, the AI-disclosure tax, and the responses-to-reviewers tax have all grown in the same period. None existed in its current form in the 1980s. Collectively, they absorb, according to section 2, around 70 million researcher-hours per year on the author side alone, with a comparable reviewer-side cost.

[25] Jones, B. F. (2009), 'The burden of knowledge and the "death of the renaissance man": Is innovation getting harder?', Review of Economic Studies 76 (1): 283–317.

This does not mean that the submission apparatus is *the* cause of the productivity slowdown. The causal pattern is far too complex for that, and other contributors (the rising complexity of frontier science, the maturation of major fields, the financialisation of higher education, the proliferation of administrative compliance across the research sector) are all plausibly larger. But the apparatus is a contributor. It is a measurable contributor. And it is the contributor that should be easiest to remove, because the function it performs can, for the first time in 2026, be performed another way. That is the argument of section 5, which begins the constructive part of this article.

5. What AI could do, narrowly

Return to Figure 2. Of the fifteen to eighteen hours of formatting and interface labour per published paper, the solid black bars (body reformatting, reference reformatting, figure reformatting, web-form metadata entry, cover letter drafting, file uploads) total approximately twelve hours. Each is, with current commercially available tools, a problem AI can compress to minutes.

Body reformatting, in any journal whose style sheet is documented, is solvable. A model with the source manuscript and the target style sheet can produce a journal-formatted document in seconds. Reference reformatting is easier still: a reference list expressed as structured data can be transformed among Chicago, APA, Vancouver, Springer author-date, and Nature numbered styles by a Citation Style Language processor that has existed in mature form since 2010.[26] Figure reformatting is slightly harder, because resolution and file-type constraints vary, but well within the capacity of any current image-processing pipeline. Web-form metadata entry can be eliminated entirely if the journal accepts the manuscript file as the source of metadata, which most journals could, but only a few do. Cover letter drafting from the abstract and contribution paragraph takes a model thirty seconds, with the author retaining a vetting step for high-stakes venues. File uploads are a matter of agentic browser automation, which is an early-commodity capability in 2026 and will be mature within two years.

[26] The Citation Style Language (CSL) specification was developed from 2003 onward and reached its current stable form around 2010. It is the basis of reference handling in Zotero, Mendeley, Pandoc, and other scholarly writing tools, and provides machine-readable style definitions for the great majority of journals indexed in the major databases.

The diagonally hatched bars (suggested reviewers and conflict-of-interest checks, funding and ethics statements, residual friction) total approximately four to five hours per published paper. These cannot be compressed to minutes, because they require human judgement that AI does not yet have and may never have. The author must vet suggested reviewers for genuine conflicts the AI cannot detect: the postdoc she fell out with, the colleague whose sister she taught, the rival who attacked her last paper. The author must authorise institutional declarations that have legal weight binding her and her institution. The author must read the consent screens. None of this is going below forty-five to sixty minutes for a paper the author is willing to put her name to.

So here is the headline estimate. If current commercial tools (Paperpal, Resub, SciSpace, and a handful of competitors) were universally adopted, the per-paper burden would fall from 15 to 18 hours to roughly 2 to 3 hours, an 80 to 85 per cent compression.[27] If autonomous browser-driving agents handle the interface end-to-end with author approval at the close, the per-paper burden falls to 45 to 90 minutes: a 90 to 95 per cent compression. Applied to the central global figure of around 70 million hours per year, the absolute savings are on the order of 55 to 65 million hours per year, equivalent to 27,000 to 32,000 person-years of researcher time recovered annually. That is the labour-equivalent of a major scientific collaboration, donated back to the activity the collaboration was designed to support. Admittedly, there are two complications.

The first is that AI introduces new overhead. Many journals now require disclosure statements about AI use, several with non-trivial documentation requirements; publisher author guidelines have grown substantially to handle this. Publishers' own AI-screening systems generate false-positive flags that authors must contest. Paper-mill detection systems, themselves AI, sometimes catch legitimate submissions and require human appeal. The reformatting tax risks being replaced by an attestation tax. My estimate is that this eats up 15 to 20 per cent of the projected savings, leaving the realistic recovery at perhaps 22,000 to 27,000 person-years per

[27] Commercial AI-assisted submission tools: Paperpal (paperpal.com), Resub, SciSpace. Capabilities vary; the 80-to-85-per-cent compression figure is the high-end estimate under universal adoption, not a current average.

year. Still substantial: equivalent, every year, to the dedicated effort that produced the Higgs boson. But the number deserves some honest discounting.

The second complication is more philosophical and points to sections 6 and 7. The narrow technical compression, using AI to do the work humans used to do, leaves the apparatus intact. It just reduces the labour of feeding it. This is a partial fix, and a fragile one, because it preserves the institutional incentive to add more layers. If publishers can extract fewer hours from authors per paper, there is a real risk they will be tempted (often for good reasons of compliance, governance, and audit) to require more documentation, more disclosure, more attestation, until the new equilibrium is again fifteen hours per paper and the apparatus is twice as elaborate as before. Anyone who has watched the development of expense-report software, tax-compliance software, or healthcare administration software in the last twenty years knows how this dynamic works. The technology compresses the per-task labour; the institution responds by increasing the number of tasks. The total labour cost stays roughly constant.

Narrow AI compression is therefore necessary but insufficient. The deeper question is whether the apparatus needs to exist in its current form at all. AI is not just a tool to do existing work faster; it is the technology that reveals how much of that work we could cease to require. To see this, we need to imagine what a journal architecture would look like if built today, from first principles, using the technology now available. That is the exercise of section 6.

6. The AI-born journal

Let me start by introducing a distinction. I shall call a journal that retrofits AI tools onto an unchanged (e.g., Editorial Manager) workflow *AI-assisted*. It does the old work faster. Instead, a journal that was designed, in 2026, around what AI agents can reliably do, and that never acquired the procedural layers we now inherit, can be called *AI-born*. It is not faster at the old work; it is differently constituted. The distinction matters because the first kind leaves the incentive structure intact (which, as section 5 warned, pushes back toward the same equilibrium), whereas the second does not. AI-born journals are, in principle, permanently lighter.

Even if not exactly in these terms, the distinction has been challenged on the grounds that incumbent systems are no longer purely manual. Editorial Manager has shipped Xtract, an automated metadata-extraction feature, since version 14.1 in 2017, with subsequent releases extending the fields that can be pulled from author-submitted Word files. The extraction populates title, authors, affiliations, and abstract, reducing the per-author re-keying burden by a measurable but modest amount. ScholarOne offers comparable features. The objection runs: are these systems not already moving toward what I call AI-born? They are not. They are doing exactly what the AI-assisted category names: bolting on automated capabilities inside a workflow whose underlying architecture—the form-filling sequence, the metadata schemata, the format requirements, the staged decision points—was designed for a manual era and has not been redesigned around what AI can now do. Xtract is a feature; an AI-born journal is an architectural commitment. The two are not on the same axis. To put the matter sharply: a 2017 metadata-extraction tool grafted onto a 2001 submission workflow is not the same as a 2026 workflow built around the assumption that metadata extraction, integrity screening, reviewer matching, and decision-letter synthesis are AI primitives the architecture is allowed to depend on. The bolt-on approach yields incremental savings inside an unchanged system; the architectural approach yields a different system. The existence of bolt-on features in incumbent systems is, if anything, evidence for the argument: the components are mature enough to bolt on, which means they are mature enough to design around.[28]

It is worth being explicit about what makes a journal AI-born rather than AI-assisted. Five necessary features distinguish the category. *First*, author-native input: the submission accepts the manuscript in whatever format the author wrote it, with no required reformatting at any stage. *Second*, manuscript-as-source: the manuscript file is the canonical source of metadata, not a separate web form to be re-entered. *Third*, agentic procedural workflow: integrity checks, routing, reviewer discovery, and production are performed by AI agents rather than by editorial-office staff or by author-completed forms, with editor override at every decision point. *Fourth*, an

[28] For a systematic survey of editorial-practice innovations across publishers, including the bolt-on category I describe, see Horbach, S. P. J. M. & Halffman, W. (2020), 'Innovating editorial practices: academic publishers at work', Research Integrity and Peer Review 5: 11.

explicit procedural-versus-substantive boundary: the architecture distinguishes where AI is operating procedurally (extraction, screening, matching, drafting) from where humans are operating substantively (editorial synthesis, peer judgement, editorial signature), and treats the boundary as a design commitment rather than a transient implementation detail. *Fifth*, multi-clock evaluation: success is measured against author time, editor time, reviewer time, queue length, and integrity outcomes, not against any single metric. A sixth feature is sufficient but not necessary: format-free output, in which the published paper can be rendered in any format a reader requests, is an architectural advantage but not a definitional requirement. A journal that meets all five necessary features is AI-born; a journal that adds AI tools to its existing workflow without meeting them is AI-assisted. Most existing pilots demonstrate one or two features; none yet demonstrates all five at scale.

6.1 The architecture

Here is a sketch of one such architecture. It is a sketch, not a specification; details will vary by field and by publisher. But the architecture's shape follows from the technology and is robust across those variations. I anchor the sketch to the three clocks introduced in §2, because the architecture compresses each of them by different mechanisms and the rhetorical force of the proposal depends on seeing this. The first agent (metadata extraction) compresses the first clock, by absorbing the per-author formatting and form-filling labour. The second and third agents (integrity checks, routing) compress both the second and third clocks, by replacing slow editorial-office processing with near-instant automated screening. The fourth agent (reviewer matching) compresses the second and third clocks together, by collapsing the editor's reviewer-search time and the elapsed days the paper spends waiting for reviewer assignment. The fifth agent (decision-letter synthesis) compresses the second and third clocks, by reducing the editor's drafting time and the wait between final review and decision. The sixth and seventh stages (production, format-free output) compress the third clock, by reducing post-acceptance time from weeks to hours. The compression is asymmetric across the clocks, but every clock is touched, and the substantive editorial signature is left untouched throughout.

The author posts the manuscript once; in whatever format she wrote it. She does not reformat for the journal; the journal accepts what she sends. An agent embedded in the submission system extracts metadata from the manuscript file itself: title, authors, affiliations, abstract, keywords, references, and funding sources. The author confirms what the agent has extracted, corrects any errors, and signs off. This step takes ten minutes.

A second agent runs the integrity checks. Image manipulation detection (Bik-style forensics now available off the shelf), paper-mill fingerprint matching against the STM Integrity Hub, statistical sanity (statcheck and GRIM on reported numbers), reference verification against Crossref and OpenAlex, conflict-of-interest scanning against the author's Scopus collaboration record.[29] The integrity checks are non-blocking: they produce a report for the editor, not an automated rejection. The author sees the report and responds if needed. A point worth stressing: integrity checks are not purely administrative. In some fields—statistics-heavy ones in particular—a statcheck or GRIM flag is itself a partial epistemic intervention into the paper's content, and the architecture should treat such checks as evidence the editor weighs rather than as automated filters. The procedural-versus-substantive boundary I named above runs through these tools as well: image manipulation detection and paper-mill matching are screening; statistical sanity and reference verification, in fields where they are decisive, are evaluation. The architecture handles screening automatically; evaluation feeds into the editorial signature.

A third agent handles routing. If the journal is part of a consortium or a cascade, the agent identifies the most likely fit among participating venues; if the journal stands alone, the agent screens for scope and recommends desk-rejection or onward review. The editor confirms or overrides. This step takes the editor perhaps fifteen minutes per submission.

A fourth agent identifies plausible reviewers. From the bibliography, from the field's recent publication record, from co-citation patterns, from public expert

[29]The Bik image-forensics methodology was systematised by Elisabeth Bik in her work on image-duplication detection in biomedical papers; statcheck was developed by Michèle Nuijten and colleagues at Tilburg; GRIM (Granularity-Related Inconsistency of Means) by Nick Brown and James Heathers. statcheck and GRIM exist as open-source tools that can run unattended on submitted manuscripts.

directories like ORCID and OpenAlex, the agent generates a candidate list of fifteen to twenty potential reviewers, ranked by topical fit, institutional independence from the author, and recent reviewing load. The editor selects three or four; the agent drafts and sends invitations and manages the reminder cascade. The editor's time on this step compresses from perhaps two hours per paper to fifteen minutes.

The reviewer who accepts the invitation receives the manuscript bundled with a structural pre-analysis: claims extracted, methodological summary, statistical audit, comparison with cited literature, list of cited papers the reviewer has not herself worked on. The pre-analysis is not a substitute for the review. It is scaffolding that lets the reviewer skip the orientation and get straight to the substantive question, which is the only question the reviewer is uniquely qualified to answer: is this contribution real, is the framing right, does the field need it? The reviewer's orientation time falls; her substantive engagement time may rise where the pre-analysis surfaces issues she would otherwise have missed. The realistic expectation is that total reviewer time per paper falls modestly, not dramatically, and that the gain is in the quality of the substantive engagement rather than in raw hours saved.

When the reviews come in, a fifth agent synthesises them into a draft decision letter for the editor. The synthesis identifies points of agreement and disagreement, flags reviewer bias if detectable, drafts a recommendation, and prepares boilerplate response language. The editor reads the reviews directly, modifies or overturns the draft recommendation, and writes the substantive paragraphs that bear her signature. The editor's time on the decision letter compresses from perhaps three hours to one and a half.

Accepted papers go through automated copyediting and typesetting (already a near-solved problem at most major publishers) and come out the other end in any format the reader wants: a long PDF, a structured-data record, a plain-language summary, an audio version, a translation. The author sees the final version once and signs off.

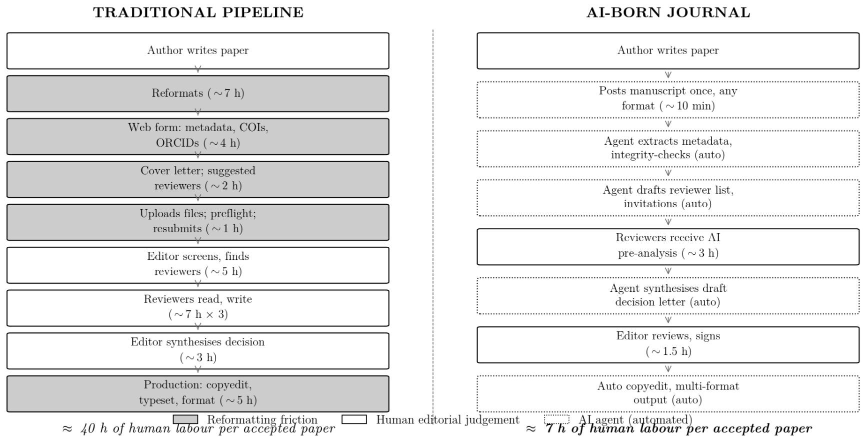

Two architectures, side by side.

Figure 6. Two architectures, side by side. Shaded boxes on the left are reformatting friction in the traditional pipeline; dotted boxes on the right are handled by AI agents; solid-bordered boxes in both columns require human editorial or reviewer judgement.

Total author time across the whole process: forty to sixty minutes, plus whatever revision is required after the reviews. Total editor time: ninety minutes per paper of substantive judgement. Total reviewer time: perhaps three hours per reviewer in straightforward cases, with orientation reduced and substantive engagement preserved or increased where the pre-analysis surfaces deeper problems. The compression operates principally on the procedural and orientation layers; substantive review time depends on the paper. Across all parties, an accepted paper in a typical case consumes roughly seven hours of human labour on the publication apparatus, against the forty hours of the traditional pipeline, with the caveat that the compression is concentrated in the procedural components and not uniform across them.

6.2 Governance

The architecture as sketched will fail without an explicit governance and risk layer, which is part of the proposal, not an implementation detail. Seven risks are constitutive enough to demand named controls in any AI-born journal. *First*, manuscript confidentiality: AI agents that process unpublished manuscripts must operate under contractual prohibitions on retention beyond the editorial cycle and on any use of submitted text for model training, with audit-loggable enforcement. *Second*, reviewer anonymity: agents that handle reviewer identities (the reviewer-discovery agent and the synthesis agent) must operate under information-flow restrictions that prevent leakage of reviewer identity into systems that authors or third parties can access. *Third*, automation bias: editor-facing summaries (the synthesis agent's draft decision letter, the reviewer-discovery agent's ranked list, the integrity agent's flag report) shape what the editor sees and how she sees it; the architecture must include explicit dissenting-view surfacing and randomised audit by senior editors to detect drift toward the agent's recommendations. *Fourth*, false-positive integrity flags: the integrity agent will sometimes flag legitimate work; the architecture must include a clear appeal path and a non-blocking review by a human integrity officer before any flag becomes a rejection. *Fifth*, reviewer-matching bias: agents that recommend reviewers based on co-citation and publication patterns will reproduce the field's historical biases; the architecture must include diversity criteria that the editor can set, and the agent must satisfy alongside topical fit. *Sixth*, human override and audit: every agent decision must be reversible by the editor without procedural friction, and every override must be logged in a record the journal's editorial board can inspect on a rolling basis. *Seventh*, adversarial-input and tool-use security: submitted manuscripts, supplementary files, and reviewer correspondence must be treated as untrusted inputs, with prompt-injection containment for documents that may contain instructions targeting the agents, sandboxed file processing, model-version logging that records which model handled which screening task, draft recommendation, or agent output, and explicit prohibition on reviewers uploading confidential manuscripts to unapproved external AI systems. The point of naming these seven controls is that procedural infrastructure is not epistemically neutral. Routing, reviewer matching, integrity flags, and reviewer-comment summaries all influence what the editor sees and how she sees it, and an architecture that says 'the editor retains judgement' without naming the controls that protect that judgement from being shaped by the agents around it is not yet a serious proposal. The AI-born journal accepts this constraint as part of its definition: the five necessary features above are constitutive of the architecture, the seven controls just named are constitutive of its legitimacy.

This is not science fiction. The pieces exist, scattered across pilots that each demonstrate a different aspect of the architecture. aiXiv, developed in 2025 by researchers at the University of Toronto, Westlake University, and the University of Electronic Science and Technology of China, demonstrates the AI-reviewer pipeline at work: five high-performing AI models review each paper, majority acceptance triggers posting, and the cycle runs rapidly compared to conventional review.[30] *eLife*, which announced its review-then-publish model in October 2022 and implemented it from 2023, demonstrates the removal of the binary accept/reject decision, posting every reviewed paper alongside its reviews.[31] Octopus, launched by Alexandra Freeman, demonstrates the decomposition of the publication unit into eight independently reviewable stages. Manubot demonstrates a GitHub-native collaborative writing and publishing pipeline, with automated citation handling, format conversion, and continuous deployment.[32] None of these, taken alone, is the AI-born journal I have sketched. None has the prestige to draw the best papers in any major field. But each shows a working component, and taken together, they show that the components already exist in pilot form somewhere in the world.

What is missing is not technology but a respected venue with established editorial judgement that has put the components together and made the integration into a public commitment. This is the move I have argued elsewhere—in a longer treatment of authorship and answerability under conditions of LLM use—should be made within a four-tier governance framework that distinguishes substantive AI use from procedural AI use.[33] In that framework, the submission apparatus is the procedural floor: the part AI can absorb almost entirely without disturbing the substantive

[30]Zhang, Pengsong, et al. (2025), 'aiXiv: A Next-Generation Open Access Ecosystem for Scientific Discovery Generated by AI Scientists', arXiv preprint 2508.15126.

[31]eLife's review-then-publish model was announced in October 2022 and implemented from 2023. See the eLife editorial from 20 October 2022, 'eLife's new model: changing the way you share your research'.

[32]Octopus (octopus.ac) was launched by Alexandra Freeman of the Winton Centre for Risk and Evidence Communication at the University of Cambridge; it decomposes the publication unit into eight independently reviewable types: research problem, rationale/hypothesis, methods, results/sources, analysis, interpretation, applications/implications, and peer review. Manubot (manubot.org) is an open-source system for continuous collaborative writing and publishing on GitHub.

[33]The four-tier framework is developed in Floridi, L. (forthcoming), 'AI and the future of publishing: authorship, answerability, and the four-tier governance of large language models in scholarly work', manuscript under review. Briefly, the framework distinguishes four modes of LLM involvement in scholarly work: (1) procedural assistance (formatting, translation, search), (2) analytic assistance (literature mapping, statistical checks), (3) generative assistance (drafting, summarisation, restructuring), and (4) substantive contribution (claims, arguments, novel analysis). Tiers 1 and 2 do not require disclosure beyond a generic acknowledgement; tiers 3 and 4 require specific disclosure under the answerability principle the framework develops.

judgements on which the journal's prestige is built. To absorb the procedural is not to dilute the substantive. It is to free the substantive to do what it should always have been doing.

7. What a journal is for

The philosophical claim that motivates everything in this article can be stated in one sentence: a journal's distinctive contribution to the scholarly-communication system is the editorial signature on a credible claim, and everything else can, in principle, be replaced or reorganised without touching that contribution. This sentence is worth examining carefully, because the proposal in section 9 stands or falls on whether it is correct.[34]

The standard taxonomy of scholarly-communication functions, originating with Roosendaal and Geurts (1997) and extended by Van de Sompel and colleagues (2004), identifies four or five: *registration* (assigning precedence to a claim), *certification* (establishing the validity of the claim), *awareness* (disseminating the claim to those who need it), *archiving* (preserving the claim over time), and on Sompel et al.'s extension, *reward* (allocating the credit on which careers and funding depend). I take this taxonomy as the right starting point. But it is a taxonomy of functions performed by the scholarly-communication system as a whole, not by the journal as a particular institution within it. arXiv performs registration and awareness without certification. Institutional repositories perform archiving. Citation indices and search platforms perform awareness. Funding agencies and tenure committees perform reward. Each of the five functions has many possible instantiations, and the journal is one instantiation of one function. Among the five, the journal's distinctive contribution—the function that, when stripped away, leaves no other institution that performs it the

[34]Roosendaal, H. E., & Geurts, P. A. T. M. (1997), 'Forces and functions in scientific communication: an analysis of their interplay', paper presented at the Conference on Co-operative Research in Information Systems in Physics, University of Oldenburg, 1–3 September. Van de Sompel, H., Payette, S., Erickson, J., Lagoze, C., & Warner, S. (2004), 'Rethinking scholarly communication: building the system that scholars deserve', D-Lib Magazine 10 (9). The Roosendaal–Geurts taxonomy identifies four functions; Sompel and colleagues add reward as a fifth. The argument here does not depend on whether the right number is four or five; what matters is that the journal's distinctive contribution among them is certification, and that certification is performed by the editorial signature.

same way—is *certification*. The other four functions are performed by the journal but not uniquely by the journal; only certification is the journal's own.

So, the question becomes: what makes a journal's certification distinctive? Here, the editorial-signature claim does its work. Certification by a journal is not the same as certification by an institutional review board, by a regulatory agency, or by a community vote on a preprint. It is certification by a named group of editors (henceforth I shall simplify and speak of a single editor) which has accepted personal accountability for a published claim, whose reputation is invested in that claim being credible, and who has positioned that claim within a community of inquiry whose standards she helps to set. Four sources of normative force converge in the editor's signature, and it is worth distinguishing them.

First, *answerability*: the editor's acceptance of the paper is institutionally traceable to her, whether or not her name appears as a byline, and a publication that turns out to be wrong—falsified, fraudulent, or methodologically broken—reaches her reputation directly through the journal's accountability chain.

Second, *situated judgement*: the editor evaluates the claim from inside a community whose standards, controversies, and live disputes she has internalised through years of work; her judgement is calibrated, not algorithmic.

Third, *institutional trust*: the journal as an institution accumulates credibility over decades, and the editor inherits and contributes to that accumulation; readers trust her judgement partly because they trust the institution she serves.

Fourth, *responsibility under uncertainty*: editorial decisions are made in conditions of evidential incompleteness (a paper is neither obviously right nor obviously wrong, and the editor must call it), and the willingness to accept responsibility for such calls is constitutive of the role.

Each of these four sources can, in principle, be borrowed by other institutions. What is distinctive is that the editor's signature concentrates all four in a single accountable act: a named person, embedded in a community, backed by an institution, taking responsibility for a claim under uncertainty. No preprint server, content-moderation platform, repository, or AI agent does all four at once. That is why the editor's signature is the journal's distinctive contribution, and why, when the signature is stripped away, what remains is no longer a journal in the sense the word has meant since 1665.

One may object that peer review is partly *constitutive* of the editor's signature, rather than merely a mechanism that supports it. The objection is reasonable. In many fields, an editor's decision is trusted because it is embedded in a peer-review institution, not merely because the editor has personal standing. Reviewers identify defects the editor would not have seen alone, raise objections from sub-specialisms outside the editor's expertise, and constrain the editor's judgement against arbitrary preference. The signature, on this view, is not the editor's alone but the editor's *ratified by* the peer-review process. I accept the objection in the following form: peer review is constitutive of the legitimacy of the signature in fields where the editor cannot plausibly be the expert on every paper she handles, which is most fields most of the time. But this concession does not deflate the editorial-signature claim; it sharpens it. The editor's signature is not a solitary judgement; it is a judgement made on the basis of, and accountable to, a peer-review process that the editor convenes, evaluates, and synthesises. What is irreducibly hers is the act of accountable synthesis: weighing what the reviewers said against what the field knows, deciding which of their objections are decisive, and accepting personal responsibility for the conclusion. AI can administer and scaffold the peer-review process; it cannot perform the synthesis, because the synthesis requires the four sources of normative force just identified, and at least three of them (answerability, situated judgement, and responsibility under uncertainty) require a human bearer. The editorial signature, properly understood, is the editor's synthetic judgement on a peer-reviewed paper, with all four sources of normative force in play.

With the editorial signature so understood, the rest of the argument follows. AI is the first technology since the emergence of the modern journal apparatus that can absorb the infrastructure without touching the signature. This is not a coincidence. It follows from the specific asymmetry between what AI can and cannot do. AI is good, and improving rapidly, at the procedural and informational work that supports the signature: formatting, metadata extraction, reference matching, integrity forensics, reviewer discovery, reviewer-comment summarisation, and decision-letter drafting. It is bad, and may remain bad for reasons that have nothing to do with capability, at the four sources of normative force that constitute the signature itself: it has nothing reputational to lose, no situated standing in a community, no institutional credibility it has earned, and no genuine responsibility to bear under uncertainty. The asymmetry between the agent and the editor is the same as the asymmetry between what can be automated in a journal's work and what cannot. It is also the asymmetry between procedural infrastructure and substantive judgement, and that asymmetry is what the AI-born architecture is built around.

This is why the proposal in section 9 is for AI-born journals rather than AI-only journals. The point of the AI-born architecture is not to replace editorial judgement with machine judgement. It is to strip away the accumulated procedural infrastructure that has, over a quarter of a century, begun to obscure the signature, to the point that authors now spend fifteen hours navigating a system, and a reader examining a published paper sees more metadata than editorial statement. The reader should see the editor's judgement first, and the infrastructure second. AI-born journals make that priority visible again by compressing the infrastructure to the point where it no longer competes for the reader's (and the author's) attention.

A philosopher or historian of science may reply that the comparison between Oldenburg's 1665 function and the contemporary editorial function is too clean. Oldenburg vetted personal correspondents and printed their letters; he was not managing the international peer review of millions of papers a year. The modern editorial function has acquired scale-dependent properties (the procedural rigour that makes a rejection defensible at scale, the institutional standing that lets an editor's 'no' carry weight in a tenure case, the editorial boards that are correspondence networks scaled into committees) that Oldenburg's function lacked. This objection, too, is well taken. But the core of the function remains, surviving when the infrastructure is stripped away. A person with reputational standing, embedded in a community, backed by an institution, accepts responsibility for a claim. Infrastructure scales the function. AI replaces the infrastructure. The core persists. That is why the restoration that the AI-born architecture offers is not a nostalgic return but a structural simplification: the journal becomes again what it was, at a larger scale, with less friction.

Admittedly, the argument depends on the substantive editorial signature being something AI cannot replace (that the synthesis just described is irreducibly human), whereas the formatting of a manuscript or the matching of reviewers is not. If, in the next decade, this argument turns out to be wrong, and substantive editorial synthesis also yields to AI compression, then the AI-born architecture collapses into pure automation, and the journal becomes a content-moderation engine after all. I do not believe this, for the reasons just given: at least three of the four sources of normative force—answerability, situated judgement, and responsibility under uncertainty—require a human bearer; institutional trust, the fourth, is the institutional medium through which those human acts acquire durable force. But I am happy to acknowledge these conditions of falsification.

The institutional implications are worth spelling out because they explain why the proposal I present in section 9 addresses publishers as partners rather than as obstacles. If the procedural infrastructure is no longer necessary to perform the act of certification, then the journal brand (which was historically a proxy for the infrastructure, the editor, and the back catalogue together) must be re-evaluated. What does *Nature* sell, in 2030, when the infrastructure is automatable? It sells the editorial judgement of its editors, the reputation of its back catalogue, and the network effect of its readership. All three are real and durable. The infrastructure, which some have historically treated as part of the brand, turns out to have been incidental. The brand does not weaken when the infrastructure is automated; if anything, it strengthens because the signature becomes more visible when it is not buried under the procedural layers that once surrounded it. The same analysis applies to *Philosophy & Technology*, to every Springer-Nature title, and to every journal anywhere whose authors currently spend fifteen hours per paper filling in web forms. The editorial judgement is the asset; the infrastructure is the cost. Any architecture that reduces costs without touching the asset is an opportunity for the publisher, not a threat. That is the argument of section 8.

8. Why this round is different

Scholarly publishing has encountered cost-compressing technology before. It will help the proposal to be explicit about that history, because the past is also what makes the future unusual.

In the 1980s, digital typesetting compressed the production cost of journals by an order of magnitude. The compression happened; subscription prices, in real terms, did not fall; and the publishing industry used the savings to increase profits, absorb rising fixed costs, and invest in new capabilities (electronic archives, online distribution infrastructure, metadata services) that authors and readers eventually came to expect.

In the 1990s, the World Wide Web compressed the distribution cost of journals to near zero. In this case, too, the compression happened. Library subscription bills, in real terms, rose. The publishing industry used the savings to increase profits, build global electronic-access platforms, consolidate acquisitions into bundles that simplified licensing for libraries (the 'big deal'), and to fund the services (reference linking, persistent identifiers, full-text search) that now underpin scholarly communication.

In the 2000s, the open access movement shifted the cost from readers to authors through article processing charges. The shift happened. APC-funded journals became a substantial portfolio, estimated by Delta Think at just under $2.4 billion for the open-access segment of the scholarly journal market in 2024.[35] The shift achieved its original goal of making published research freely available to readers. It also, as sceptics predicted, came with unintended consequences: APC pricing at the premium end (*Nature Communications*'s APC is now $7,350) has translated into a substantial author-side cost burden that the original advocates of OA did not anticipate.[36] The *PLOS ONE* APC sits at $2,477, a level that has made megajournal publication very expensive for much of the world.[37]

In each of these rounds, the technology compressed a cost that the publisher bore, and the saving accrued, in the first instance, to the publisher. This was not a conspiracy; it was simply where the cost was, and therefore where the savings landed. In each round, researchers benefited too, but indirectly and with a lag, and often less than the first-order analysis predicted.

[35]Delta Think (2025), "Market Sizing Update 2025: Has OA Recovered Its Mojo?" The $2.4 billion figure is Delta Think's estimate for the open access segment of the scholarly journal market in 2024, reported in the 2025 update.

[36]The Nature Communications article processing charge was $7,350 as of early 2026. For current pricing, see nature.com/ncomms.

[37]The PLOS ONE APC sits at $2,477 for most articles as of 2026, with partial waivers and discounts available for authors from lower-income countries. See plos.org/publish/fees.

This round is different in a specific and important way. The cost being compressed is not borne by the publishers but by the researchers. The savings, if realised, accrue in the first instance to the researchers, in the form of fifteen hours per paper returned. This is the first round of publishing-technology compression in which the incentive structure and the distributional consequences align: researchers want the change, publishers can implement it, and neither side must wait on the other to benefit.

For the publisher, the AI-born architecture offers three first-order advantages. It reduces the operational costs of the submission and production pipeline, which are rising under pressure from paper mills, AI-generated submissions, and growing compliance requirements. It shifts the publisher's competitive differentiation away from the apparatus (which is costly, becoming commoditised, and which no publisher now wants to compete on) toward editorial judgement and brand reputation (which are defensible and which the major publishers already own). It provides, for the publisher willing to move first, a strong public narrative at a time when the industry is otherwise under pressure to explain its value. These are not small benefits. They align, cleanly and for the first time in decades, with the benefits to the researchers.

There is a further consideration, which is obvious and should not be perceived as adversarial. The AI-born architecture I have sketched does not require Springer-Nature or any incumbent to act. Every component—the integrity-checking agent, the reviewer-matching agent, the metadata extractor, the decision-letter synthesiser—is publicly available or within close reach. A well-funded startup, a consortium of learned societies, or one of the major AI laboratories that has expressed interest in scientific infrastructure could build a credible AI-born journal in a period measured in months rather than years. Several are, at the time of writing, already building pieces of it. The pattern is familiar from adjacent industries. An incumbent operates behind what appears to be a durable moat; a new technology arrives that makes the moat's underlying function commoditisable; the moat collapses faster than anyone expected; and the incumbents who assumed the moat was permanent discover that it was an accident of the pre-disruption era rather than a source of durable value. I do not think this is a distant prospect for scholarly publishing. I think it is the likely trajectory of the next five to ten years. Traditional publishers could easily be 'uberised'.[38]

The question, then, is not whether AI-born journals will exist. They will. The question is whether the venues that carry them will be launched from inside the established scholarly-publishing ecosystem, with the editorial boards, the back catalogues, the brand trust, and the community standing that they have built over decades, or whether they will be launched from outside, by actors who have the technology but not the reputational capital, and who will spend the next decade accumulating that capital at the established publishers' expense. The first outcome is better for everyone, including the incumbent: the community keeps the editorial institutions it trusts, the transition is faster and less disruptive, and the publishers capture the upside rather than watching it accrue elsewhere. The second outcome is worse for everyone, but it is fully within reach of actors who do not need the incumbent's permission to proceed. This is what I mean when I say the incentive alignment is now genuine. Not that the publishers should do this because it is nice, but because the alternative—declining to do this—is not a return to the status quo. The status quo is ending, with or without the traditional publishers' participation. The only choice is whether they are the ones who will lead the transition or be left behind by it.

The point is worth emphasising because it changes the politics. In previous rounds, the researcher community and the publisher had to negotiate over how a newly compressed cost would be divided. In this round, there is no division to negotiate; the cost being compressed was never the publisher's to begin with. What the publisher

[38]The reference case is the New York City taxi medallion, which traded for approximately $1.3 million in 2013 and had collapsed to under $200,000 by 2019 following the arrival of ride-hailing platforms. The analogy is imperfect—scholarly journal brands are not regulated quantities like medallions, and editorial judgement is a genuinely durable asset in a way that taxi dispatching was not—but the structural dynamic is instructive. In both cases, an incumbent's apparent moat turned out, on the arrival of the disrupting technology, to have been an accident of the pre-disruption era rather than a source of value. The New York medallion was valuable when arranging a taxi required physical dispatch infrastructure only medallion-holders could afford to operate; once the smartphone made arranging a taxi trivial for anyone, the medallion was left protecting nothing. The submission apparatus, the in-house editorial office, the production pipeline, and the distribution platform were, similarly, the scarcity the scholarly publisher traded on. The AI-born architecture removes the scarcity. What remains—editorial judgement, brand, back catalogue, community—is real and defensible, but it is a different asset class and competes on different terms. The publishers that recognise this early will capture the transition; those that do not will discover that the moat they were defending was not the moat that mattered.

gains is the opportunity to deliver the compression, with all the reputational credit that comes with being first. This is what I mean by a partnership opportunity.

9. Conclusion: a proposal

I write this article as the editor of *Philosophy & Technology*, a Springer-Nature journal whose submission workflow runs on Editorial Manager. I am also aware that, as editor, I am part of the system I describe. Every Editorial Manager page the authors navigate was approved at some point by the editorial team I lead. None of the points above can be made sincerely without this acknowledgement. So, what I propose, specifically, is the following. I would like to work to develop *Philosophy & Technology* into a pilot AI-born journal over the next twenty-four months. Because the AI-born architecture is defined partly by multi-clock evaluation, the pilot could be measured against the criteria the architecture itself sets. I am not proposing a revolution but a controlled experiment at a single willing venue, with the results to be reported, in due course, in this journal. If the experiment succeeds, *Philosophy & Technology* becomes an example of an AI-born journal in the humanities, and Springer-Nature becomes a publisher that leads the move rather than follows it. If the experiment reveals unexpected difficulties, those difficulties are themselves valuable findings, and the research community learns from them. The cost of the experiment, in real terms, is small. The potential value, for the publisher and the community, is large.

A final reminder, to be utterly clear. I am not proposing to replace human peer review with AI review, or to have editorial decisions made by algorithms. I am not proposing to weaken the disclosure requirements that protect scholarly integrity. The AI-born architecture absorbs the procedural layer, not the substantive one. Reviewers still review; editors still decide; the signature still signs. What should change is how much of the work surrounding those irreducible acts is performed by the people who used to perform it, and how much of it is handed over to artificial agents that perform it more quickly, more consistently, and, where the task is well-defined, more accurately.

A billion hours is not a small number. Half a million person-years is not a small number. We have spent them. We can stop spending them. The decision is not technical; the components are sufficiently ready for serious pilots, even if the integrated governance, procurement, and audit infrastructure must still be proven at the journal scale. The decision is institutional. And institutions, as Oldenburg knew on that Monday in March 1665 when he posted sixteen pages of news from the learned world, are made by people who decide to act.

References

Aad, G., et al. (ATLAS and CMS Collaborations). 2015. "Combined Measurement of the Higgs Boson Mass in pp Collisions at $\sqrt{s} = 7$ and 8 TeV with the ATLAS and CMS Experiments." *Physical Review Letters* 114 (19): 191803. https://doi.org/10.1103/PhysRevLett.114.191803.

Aczel, Balázs, Barnabás Szaszi, and Alex O. Holcombe. 2021. "A Billion-Dollar Donation: Estimating the Cost of Researchers' Time Spent on Peer Review." *Research Integrity and Peer Review* 6: 14. https://doi.org/10.1186/s41073-021-00118-2.

Andersen, Mikkel Zola, Siv Fonnes, and Jacob Rosenberg. 2021. "Time from Submission to Publication Varied Widely for Biomedical Journals: A Systematic Review." *Current Medical Research and Opinion* 37 (6): 985–993. https://doi.org/10.1080/03007995.2021.1905622.

Björk, Bo-Christer, and David Solomon. 2013. "The Publishing Delay in Scholarly Peer-Reviewed Journals." *Journal of Informetrics* 7 (4): 914–923. https://doi.org/10.1016/j.joi.2013.09.001.

Bloom, Nicholas, Charles I. Jones, John Van Reenen, and Michael Webb. 2020. "Are Ideas Getting Harder to Find?" *American Economic Review* 110 (4): 1104–1144. https://doi.org/10.1257/aer.20180338.

Chen, Tzu-An, Ming-Hwai Lin, Yu-Chun Chen, and Tzeng-Ji Chen. 2024. "The Time from Submission to Publication in Primary Health Care Journals: A Cross-Sectional Study." *Publications* 12 (2): 13. https://doi.org/10.3390/publications12020013.

Delta Think. 2025. "Market Sizing Update 2025: Has OA Recovered Its Mojo?" *News & Views*. https://www.deltathink.com/news-views-market-sizing-update-2025-has-oa-recovered-its-mojo.

Dreier, Casey. 2022. "An Improved Cost Analysis of the Apollo Program." *Space Policy* 60: 101476. https://doi.org/10.1016/j.spacepol.2022.101476.

eLife. 2022. "eLife's New Model: Changing the Way You Share Your Research." Editorial, *eLife*, 20 October. https://elifesciences.org/inside-elife/54d63486/elife-s-new-model-changing-the-way-you-share-your-research.

Ellison, Glenn. 2002. "The Slowdown of the Economics Publishing Process." *Journal of Political Economy* 110 (5): 947–993. https://doi.org/10.1086/341868.

Fishman, Charles. 2019. *One Giant Leap: The Impossible Mission That Flew Us to the Moon*. New York: Simon & Schuster.

Floridi, Luciano. Forthcoming. "AI and the Future of Publishing: Authorship, Answerability, and the Four-Tier Governance of Large Language Models in Scholarly Work." Manuscript under review.

Horbach, Serge P. J. M., and Willem Halffman. 2020. "Innovating Editorial Practices: Academic Publishers at Work." *Research Integrity and Peer Review* 5: 11. https://doi.org/10.1186/s41073-020-00097-w.

Jiang, Yan, Robert Lerrigo, Anika Ullah, Muthu Alagappan, Steven M. Asch, Steven N. Goodman, and Sidhartha R. Sinha. 2019. "The High Resource Impact of Reformatting Requirements for Scientific Papers." *PLOS ONE* 14 (10): e0223976. https://doi.org/10.1371/journal.pone.0223976.

Jones, Benjamin F. 2009. "The Burden of Knowledge and the 'Death of the Renaissance Man': Is Innovation Getting Harder?" *Review of Economic Studies* 76 (1): 283–317. https://doi.org/10.1111/j.1467-937X.2008.00531.x.

Kovanis, Michail, Raphaël Porcher, Philippe Ravaud, and Ludovic Trinquart. 2016. "The Global Burden of Journal Peer Review in the Biomedical Literature: Strong Imbalance in the Collective Enterprise." *PLOS ONE* 11 (11): e0166387. https://doi.org/10.1371/journal.pone.0166387.

LeBlanc, Allana G., Joel D. Barnes, Travis J. Saunders, Mark S. Tremblay, and Jean-Philippe Chaput. 2019. "Scientific Sinkhole: The Pernicious Price of Formatting." *PLOS ONE* 14 (9): e0223116. https://doi.org/10.1371/journal.pone.0223116.

Oldenburg, Henry, ed. 1665. *Philosophical Transactions*, 1 (1). London: John Martyn and James Allestry.

Petrou, Christos. 2022. "Publishing Fast and Slow: A Review of Publishing Speed in the Last Decade." *The Scholarly Kitchen*, 8 November. https://scholarlykitchen.sspnet.org/2022/11/08/guest-post-publishing-fast-and-slow-a-review-of-publishing-speed-in-the-last-decade/.

Petrou, Christos. 2024. "Publishing Fast or Slow: How Speed Varies for Similar Journals." *The Scholarly Kitchen*, 19 March. https://scholarlykitchen.sspnet.org/2024/03/19/guest-post-publishing-fast-or-slow-how-speed-varies-for-similar-journals/.

Petrou, Christos. 2025. "How the Growth of Chinese Research Is Bringing Western Publishing to Breaking Point." *The Scholarly Kitchen*, 8 July. https://scholarlykitchen.sspnet.org/2025/07/08/guest-post-how-the-growth-of-chinese-research-is-bringing-western-publishing-to-breaking-point/.

Phillips, Kathryn A., and Danea M. Horn. 2025. "Review and Publication Times and Reporting Across Journals on Health Policy." *JAMA Network Open* 8 (5): e2512545. https://doi.org/10.1001/jamanetworkopen.2025.12545.

Roosendaal, Hans E., and Peter A. T. M. Geurts. 1997. "Forces and Functions in Scientific Communication: An Analysis of Their Interplay." Paper presented at the Conference on Co-operative Research in Information Systems in Physics, University of Oldenburg, 1–3 September.

Tennant, Jonathan P., Jonathan M. Dugan, Daniel Graziotin, et al. 2017. "A Multi-Disciplinary Perspective on Emergent and Future Innovations in Peer Review." *F1000Research* 6: 1151. https://doi.org/10.12688/f1000research.12037.3.

Tripp, Simon, and Martin Grueber. 2011. *Economic Impact of the Human Genome Project.* Columbus, OH: Battelle Memorial Institute.

Van de Sompel, Herbert, Sandy Payette, John Erickson, Carl Lagoze, and Simeon Warner. 2004. "Rethinking Scholarly Communication: Building the System That Scholars Deserve." *D-Lib Magazine* 10 (9). https://doi.org/10.1045/september2004-vandesompel.

Zhang, Pengsong, et al. 2025. "aiXiv: A Next-Generation Open Access Ecosystem for Scientific Discovery Generated by AI Scientists." arXiv preprint 2508.15126.