Mathematical Life in the Age of AI

Understanding, judgment, and the freedom to inquire

Henry Shin · First published 12 October 2026 · Revised 3 November 2026

Refacte · web edition

Contents

I. What understanding makes possible **2**

1 Understanding as a purpose 2

2 A comparison worth understanding 4

3 Becoming capable 5

II. How support becomes authority **6**

4 A hearing before a verdict 6

5 What performed mastery costs 8

6 The means to continue 10

III. Giving greater capacity a shared purpose **11**

7 When the grounds for a judgment change 11

8 Credit as an intellectual record 13

9 Making participation substantial 14

10 An inquiry with somewhere to go 16

Notes and sources **18**

Publication history

This essay develops and combines *Mathematical Understanding in the Age of AI*, first published 7 October 2026, and *The Culture of Mathematical Judgment*, first published 12 October 2026. Their former addresses retain links to the corresponding sections and earlier editions. Revised 3 November 2026 to make the coordinate choices, allowable operations, and questions opened by the mathematical examples explicit.

A theorem can be settled long before it becomes intelligible to a particular reader. Its proof may be familiar, its explanation published, and its connections well understood by specialists. Yet coming to see why it works can change what that reader is able to think. The mathematical record need not change for an encounter with mathematics to matter.

AI gives this familiar distinction a new practical importance. Machines can contribute arguments, find organizing ideas, and explain difficult subjects. Greater capability creates choices about the activity it will serve. It can raise expectations for output and give more people the means to enter a theory, develop a question, and share what they come to understand.

These choices reach into the organization of mathematical life. People acquire their capacities through instruction, examples, criticism, and time. The institutions that judge their work also distribute many of the resources through which the work becomes possible. Their influence extends from the assessment of finished work to the conditions under which people become capable of producing it. As the means of inquiry change, so do these dependencies. A person may become able to continue without securing an established institution's approval; an institution may gain new ways to support contributions it previously could not examine.

What follows from greater capacity therefore depends on how support, judgment, and recognition are arranged. Mathematics supplies a particularly clear reason to care about those arrangements. Its explanations can become part of more people's intellectual lives without first becoming new discoveries. Its discoveries, in turn, can open questions whose value takes time and cooperation to unfold. A culture organized around understanding should make room for both.

I. What understanding makes possible

1 Understanding as a purpose

Gowers's *The Two Cultures of Mathematics* describes a difference of priority: whether problem solving chiefly serves understanding, or understanding chiefly serves problem solving. He stresses their overlap. He also reproduces Atiyah's account of inquiry whose direction emerges through curiosity and connections: "I have never started off with a particular goal, except the goal of understanding mathematics." That purpose can animate research, but it can also be fulfilled by learning mathematics already known to others.[1]

Poincaré compared a reader of mathematics to a spectator at chess. Recognizing that every move obeys the rules does not explain why the moves were chosen or how they form an organized whole. Similarly, checking the steps of a deduction leaves questions about the organization of an argument. His account of mathematical invention extends the point to learning: "I myself re-invent it in so far as I repeat it." Following another person's reasoning can involve constructive thought.[2][3]

Understanding develops through different accomplishments. A reader learns why a hypothesis matters, recognizes a mechanism, or adapts a construction to a changed example. De Regt's account in philosophy of science connects understanding with using an intelligible theory to construct

explanations, under standards of empirical adequacy and consistency. Mathematics has its own objects and standards, but the emphasis on capacities is useful here. Possessing an explanation and being able to work with it are different achievements.[4]

These capacities need not develop evenly. A researcher may use a theorem responsibly while relying on specialists for its proof. Another may check a deduction without seeing its most fruitful interpretation. Results also have value before their explanatory possibilities are exhausted: a counterexample closes a proposed route, and a proved implication can let another investigation proceed. Tao's account of mathematical quality recognizes several goods, including problem solving, technique, conceptual insight, unification, and exposition. No single performance measures them all.[5]

Machine strength makes this plurality easier to see in games. Superhuman Go did not end human play. An observational study of more than 5.8 million professional moves found higher AI-assessed decision quality and earlier departures from previously recorded sequences after superhuman AI became available. Players continued developing their practice with stronger play available elsewhere. In *Game Changer*, Sadler and Regan organized AlphaZero's chess games around strategic themes, turning unfamiliar machine decisions into material for human study.[6][7][8]

Benjamin Teuber's account of Go teaching identifies a smaller distinction. An earlier move can reveal a mistaken plan even when a later move incurs the greatest numerical loss. Locating the worst move and understanding the misconception answer different questions. An instructor can use stronger analysis to help a player grasp a plan the instructor could not have discovered unaided. AI may supply the explanation too. The value of learning, examining, and sharing it need not rest on an ability reserved to humans.[9]

Mathematics has no agreed score for the relative value of a theorem, an illuminating example, or a careful exposition. Originality concerns what a contribution adds to existing knowledge; significance, explanatory power, and what someone learns from it remain further questions. A proof must establish its conclusion. A claim of discovery needs accurate attribution. A scarce grant requires a reasoned allocation. Those standards apply to particular claims and decisions; they do not supply one ranking of every mathematical activity.

Dan Rockmore's comparison of a classroom with a gym directs attention to capacities acquired through participation. Grothendieck's objection to cerebral "weightlifting" concerned a different use of effort: allowing competitive standing to govern creative work. Practice can be worthwhile without becoming a competition, and competition can be worthwhile without defining the entire practice. A seminar that makes an established theory accessible has fulfilled a mathematical purpose even if no paper follows.[10][11]

Nor does every afternoon need to demonstrate that purpose. Some attempts end without a result or an identifiable lesson. Replacing a demand for constant publication with a demand for constant demonstrable insight would preserve much of the same pressure. Mathematical life can accommodate play, failed attempts, patient study, and the pleasures of comparative success. Greater capability enlarges these possibilities without deciding which ones people should pursue.

2 A comparison worth understanding

The Cayley–Hamilton theorem says that substituting a square matrix into its characteristic polynomial gives zero. To understand its cancellation, begin with a question about walks: can a repeated visit to a vertex tell us which terms to pair?

Let $A = (a_{ij})$ be an n-by-n matrix over a commutative ring with identity. Regard a_{ij} as the weight of a directed edge from vertex i to vertex j. Matrix multiplication then has a concrete reading: $(A^m)_{ij}$ sums the weights of all length-m walks from i to j, multiplying edge weights along each walk. A length-zero walk contributes one when its endpoints agree.

A length-n walk on n vertices must repeat a vertex. Cutting out the intervening closed segment leaves the same endpoints, but loses the weight of that segment. Keep it as a reserved cycle instead. The candidate correction is now a shorter walk together with the information removed from it. Several cuts may require several reserved cycles; the next task is to choose a cut that can be undone.

Write

$$\chi_A(t) = \det(tI - A) = \sum_{k=0}^{n} c_k t^{n-k}.$$

Expanding the determinant gives another kind of object. The coefficient c_k is a signed sum of collections of vertex-disjoint directed cycles occupying k vertices, including loops. Each collection contributes the product of its edge weights and one minus sign per cycle. This sign comes from combining the permutation sign with the minus signs on the selected matrix entries.

Keep a walk from i to j, together with a reserve of vertex-disjoint cycles, using n edges in total. The reserve may meet the walk. The determinant expansion shows where these candidate corrections occur:

$$(\chi_A(A))_{ij} = \sum_{k=0}^{n} c_k (A^{n-k})_{ij} :$$

the reserve supplies k edges, and the walk supplies the remaining $n - k$. We have turned a polynomial identity into a cancellation problem on finite records.

Scan a record's walk from its starting vertex. Stop at the first contact with a reserved cycle or the first repeated walk vertex. In the first case, splice the reserved cycle into the walk. In the second, cut the newly closed cycle out of the walk and put it in reserve. Before that stopping place, the visited vertices are distinct and avoid the reserve. This makes the two moves undo one another: after insertion, the first repetition closes the inserted cycle; after removal, the cycle's base becomes the first contact with the reserve.

The endpoints and the product of edge weights are unchanged. The number of reserved cycles changes by one, so the sign reverses. There cannot be a record with no stopping place. If the reserve occupies k vertices, a length-$n-k$ walk would have to visit $n-k+1$ distinct vertices among

the $n-k$ left outside the reserve. Thus every record has an opposite-sign partner, and every matrix entry sums to zero. This is Straubing's combinatorial proof, as presented by Zeilberger.[12][13]

On two vertices, the off-diagonal entry already shows the move:

$$a_{11}a_{12} + a_{12}a_{22} - a_{11}a_{12} - a_{22}a_{12} = 0.$$

The positive terms are two-edge walks; their partners are a one-edge walk with a loop in reserve. Commuting coefficients let us move that loop between the two parts without changing its weight. The complete walk-and-cycle construction works through the choices and the general pairing.

A reader can check this understanding by locating the first event in a new example, carrying out the exchange, and explaining why its partner returns to the original record. Changing the total number of edges gives a further test: below degree n, some records can remain unpaired. The same rule then identifies what survives, rather than merely reproducing a vanishing sum.

3 Becoming capable

An insight often arrives before the machinery required to explain it. Someone suspects that two descriptions should be compared, that an invariant has forgotten something important, or that a small example defeats a sweeping claim. The initial thought may be brief. Developing it requires enough background to find out what the suggestion amounts to.

A colleague might supply a reference, recognize an obstruction, or spend an afternoon exploring the unfinished thought. A text can provide examples without answering the question that brought a reader to them. AI adds another possible source of assistance, including substantive reasoning. Inspiration, calculation, and criticism can move between participants; working through details can produce the idea that later makes those details seem routine.

Bryan and collaborators describe such an exchange in their 2026 work on a Grothendieck-ring calculation. AI proofs of special cases helped the human authors recognize a general strategy; revised questions then supported AI development of complete arguments for a stronger result. Specialized internal systems and human scaffolding contributed at different stages. The authors wrote the final exposition and retained responsibility for correctness. The account exhibits reciprocal development, while leaving open how much of that arrangement a person with ordinary access could reproduce.[14]

Assistance can also change the encounter with an existing work. A reader can ask how the argument is organized, acquire missing background, test an interpretation against the text, or investigate why an example fails. A passage that would otherwise end the encounter becomes the beginning of a question. This can sustain solitary study or help someone enter a reading group with something precise to ask.

The educational effect depends on the arrangement. In a randomized experiment involving nearly a thousand high-school mathematics students, Bastani and collaborators found that two GPT-4 assistants improved performance during practice. With assistance removed, students who had used the more unrestricted interface scored 17 percent lower than the control group. A tutor designed to

support learning largely avoided that loss, without establishing an improvement in the unassisted exam. Immediate performance and acquired capacity needed separate assessment.[15]

In a randomized crossover study of two physics lessons, Kestin and collaborators found stronger immediate learning outcomes with structured AI tutoring than with active classroom instruction. The tutor incorporated expert-written solutions and careful instructional design. The result concerns those lessons, rather than an entire course of intellectual development. Together, the studies show why "using AI" describes too little to determine an educational outcome: the task, the guidance, and what the learner is asked to do matter.[16]

Some effort is an avoidable barrier; some is part of learning to see. A useful explanation can remove hours of fruitless searching while leaving a reader to reconstruct the crucial step, vary the assumptions, or explain why a tempting alternative fails. The relevant comparison is between the learning arrangements actually available. For someone without specialist help, access to examples and correction may open an inquiry that would otherwise stall.

Entry into a subject can be organized with the same care. Lawvere challenged the assumption that one inherited foundational ladder supplies the only rigorous route, while retaining genuine prerequisites and offering conceptual alternatives. A research introduction can likewise specify a local threshold: one construction, a comparison map, or a family of examples that makes the wider architecture intelligible. Identifying that threshold distributes guidance that would otherwise depend on personal access.[17]

Private study can fulfill its purpose without an audience. Organized mathematics also depends on people who can reconstruct arguments, detect errors, teach prerequisites, and maintain what others use. Acquiring these capacities often involves established mathematics; their renewal is part of what makes further discovery possible. Institutions draw on this work, and their responsibilities follow from the roles they undertake. Teaching develops participants; publication supplies judgments others rely on; funding gives some inquiries the means to continue. Their contribution to mathematics should be judged partly by how they sustain the capacities they use.

II. How support becomes authority

4 A hearing before a verdict

Doron Zeilberger once interrupted a lecture by Hyman Bass to announce that its result essentially rediscovered a classical theorem. His colleagues were impressed. Later that day he discovered that the hypotheses differed; he and Dominique Foata subsequently spent two summers establishing the connection properly. Public assurance had received its reward before the mathematical relationship had been understood. The episode identifies a danger wherever the appearance of mastery can settle whose work deserves further attention.[18]

Attention is a demanding resource. A reader may need unfamiliar definitions, years of background, or a difficult translation between frameworks. Reputation can legitimately guide the first investment of time. A record of reliable work or a knowledgeable colleague's recommendation supplies

information that a title and abstract do not. The difficulty arises when a decision about where to begin acquires the force of a verdict about what a work contains.

Galois made the distinction sharply in his unfinished preface of December 1831. A referee's inability to understand, he complained, could become a public verdict that the work meant nothing. His polemic does not establish an impartial account of every encounter. It identifies the narrower point that matters here: what an evaluator has established may be less than everything a work can offer.[19]

Three kinds of authority meet in such decisions. Epistemic authority supplies informed judgments. Administrative authority allocates scarce attention, money, and positions. Status authority confers esteem and influences whose next claim receives a hearing. The same institution can exercise all three, but success in one role does not settle the others. A paper can be outside a journal's remit, a department can need a different specialty, and a reader can lack time. These are genuine reasons to decline an investment. They become misleading when presented as demonstrations of a mathematical deficiency that was never assessed.

If a purported antecedent requires a hypothesis absent from the submitted result, a claim of redundancy needs a reduction that accounts for the difference. Identifying a defect in the claimed extension would support a different objection. A correction need not determine the publication decision. It should determine whether the claim of redundancy continues to circulate with the institution's authority.

Publishing makes the distinction visible. David Craven describes employment as his principal reason for submitting to journals, separating their career-significant endorsement from dissemination through arXiv. The arrangement of those functions can change. In 2001, forty members of *Machine Learning*'s editorial board resigned to support JMLR. Discrete Analysis later combined arXiv hosting with refereeing and public editorial introductions. Expertise could continue to serve assessment through a different organization of publication and explanation.[20][21][22]

The profession's stated commitments provide grounds for evaluating these arrangements. Merton located scientific reliability in shared knowledge, impersonal assessment, and organized scrutiny, rather than exceptional personal virtue. The European Mathematical Society asks referees to judge intrinsic merit independently of authors' standing. An institution claiming to support inquiry should be judged by how its practices serve that purpose. Demanding selection and responsibility for the grounds of a judgment belong within the same commitment.[23][24]

Status can affect the assessment itself. In a preregistered experiment at a finance journal, referees recruited through anonymous invitations evaluated the same manuscript under a Nobel laureate's name, an early-career researcher's name, or neither. About 23 percent recommended rejection under the prominent author's name, 48 percent with no name, and 65 percent under the early-career researcher's name. Attention has its own mechanisms: Simcoe and Waguespack studied busy-period announcements at the Internet Engineering Task Force that sometimes omitted authors' names, although the full drafts retained them. Omitting prominent names reduced attention and eventual publication for individual submissions. No comparable effect was detected among prescreened working-group drafts receiving more organized scrutiny; the comparison does not establish that scrutiny alone caused the difference.[25][26]

Opportunities also accumulate. Merton's Matthew-effect analysis distinguishes credit, visibility, and resources. In Bol and colleagues' threshold study of Dutch early-career funding, applicants just above an award cutoff received more than twice as much subsequent funding over eight years as those just below it, excluding the initial award. Later application behavior differed too. An initial opportunity can change the time, advice, confidence, and collaborators available for later work.[27][28]

Later achievement records what a person has done under particular conditions of support. It can provide good evidence of present competence without independently vindicating the decisions that helped produce it. Someone given time, collaborators, and room to recover from an unsuccessful attempt may become substantially more capable. That development shows what the opportunity made possible; it does not establish what someone denied the opportunity would have become. Treating the later difference as independent confirmation of superiority already present at the first decision can make an institution's allocations appear to vindicate themselves.

Institutions help form what they later assess. Their intellectual responsibility includes making competence acquired along other routes visible, and allowing an earlier refusal to be reconsidered when the work changes. A person needs a practicable way to develop the relevant capacities and return with evidence that can alter the assessment. Otherwise, the apparent stability of an intellectual hierarchy may partly reflect the difficulty of entering it.

Personal accounts show that such routes can change. Caroline Series contrasts acceptance at Berkeley, isolation at Cambridge in 1977–78, and feeling valued at Warwick, while recording later improvement. Lenore Blum describes a discouraging MIT entry encounter concerning her gender, Singer's intervention, and generally fair faculty treatment after admission. Their accounts identify decisions and conditions under which mathematical ability could find a different reception.[29][30]

5 What performed mastery costs

The most immediate cost of a status contest can be a loss of information. A seminar needs decisive criticism, but it also needs people to disclose a gap. When uncertainty is treated as inferiority, a question remains unasked because it would reveal insufficient familiarity. A prominent person's objection can settle the discussion without being examined. The room conveys information about rank while losing information about the mathematics.

Knowledge is distributed in ways that make this loss consequential. The speaker may know the construction; a listener may know a counterexample; a newcomer may notice an assumption others have ceased to articulate. Gathering capable people creates an opportunity to combine these resources. Humiliation and performed mastery make some of them harder to use. The harm lies partly in the competence that participants cannot exercise together.

Ambition itself need not produce that harm. Hardy treated intellectual curiosity, professional pride, and the pursuit of reputation as coexisting motives. May's memorial for Frank Adams combines exceptional competitiveness with kindness and care for exposition; it includes Adams seeking help for a student who shook when visiting him. Mixed motives and substantial virtues do not remove

the need to respond to the consequences of conduct. They explain why admiration alone supplies an unreliable professional safeguard.[31][32]

Grothendieck's criticism of competitive mathematics includes his own possessiveness, dismissiveness toward an unknown researcher, and preferential attention to students who enhanced his prestige. Shulman similarly includes himself when describing the egotistical impulses encouraged by mathematical status games. These accounts are useful because they place the conflict inside a love of mathematics: an attachment to understanding can become entangled with the identity of being someone who already understands. Another person's contribution then threatens an image that the inquiry ought to be free to revise.[33][34]

Reputation can protect a claim from checking as well as protect a person from embarrassment. Voevodsky described an error in his work on presheaves with transfers that survived for years before Deligne's detailed examination exposed it. He also recounted his prolonged resistance to Simpson's counterexample to work with Kapranov. His explanation connects difficult arguments and authorial reputation with inadequate scrutiny, and helped motivate his verification program. Self-correction made the account constructive: the response was to develop better means of examination.[35]

These failures can persist through ordinary dependencies. A junior researcher may need a recommendation; an editor may value a prominent person's participation; a colleague may see little advantage in pursuing an awkward objection. Where the person being challenged also controls guidance or advancement, a dismissive interruption or refusal to give reasons can carry consequences far beyond the encounter itself. Even a correct objection to an argument cannot justify penalizing a reasonable request to understand it. The risk of asking for an explanation falls on the dependent person while the authority of the judgment may survive without one. Rigor must therefore govern evaluators' claims as well as authors' arguments. Technical distinction warrants credit for an achievement, but cannot make a judgment immune to correction.

The same issue affects unfinished thought. In his June 1983 letter to Brown, Grothendieck described how disapproval of speculative thinking could inhibit its development even in private. His later Crafoord Prize refusal alleged that professional ethics had deteriorated, including appropriation of work from people unable to defend themselves. His diagnoses concern what happens when unequal recognition affects which ideas can be developed and whose account survives.[36][37]

An alternative begins in ordinary explanatory habits. If a step follows from a definition, name it. If it uses standard machinery, identify the machinery and its hypotheses. If a question reveals a missing prerequisite, give the listener a route to acquire it. Riehl criticizes language that calls arguments trivial or presumes everyone learned them in kindergarten. In the same interview she describes algebraic topology as welcoming to young women in her experience and credits deliberate Women in Topology efforts. The pairing identifies practices people can change and people who have changed them.[38]

Grothendieck's own working practice supplies a related example. Illusie recalls his insistence that even apparently trivial assertions receive a proof or reference, and Dieudonné's further development of the drafts' details and presentation. Ambitious abstraction was made usable through explicit work. Joyal's call for collective, accessible mathematical knowledge likewise accepts specialization while rejecting artificial intellectual and professional barriers. The issue is how demanding knowledge becomes available to another person.[39] [40]

Leadership can help organize that work. A strong mathematician may identify fruitful questions, secure resources, and make unusual projects possible. Shulman valued what Voevodsky's leadership had enabled, while explaining that the portrayal of homotopy type theory as one charismatic figure's program did not match his experience. Leadership earns an intellectual role through the options it creates; those it helps should remain able to acquire different interests and revise its direction.[41]

6 The means to continue

Institutions concentrate resources that are often difficult to assemble separately: instruction, criticism, collaborators, funding, and an audience. When those resources come as a bundle, the standards for entering or advancing within an institution can govern the direction of the work. A judgment about one contribution can become entangled with the power to determine whether its author has the means to continue at all.

Fifteenth-century European monasteries bought printed books and commissioned editions from secular workshops; some operated presses themselves. Johannes Trithemius, abbot of Sponheim, kept manuscripts and printed books in his library and praised printing for making learning more widely available. His defence of copying by hand, written in 1492, appeared in print in 1494. He valued copying as religious discipline, as a means of transmitting works that remained unprinted, and for the scribe's freedom to choose and correct a text independently of a printer's program.[73]

For a particular inquiry, the threshold can be local. Someone may need enough help to enter a literature, develop a reduction, test an example, or understand an objection. AI changes both what institutions can provide and what people can assemble outside them. Its effect on dependence turns on that comparison. When the available help suffices for the question at hand, access to a particular department or network becomes less decisive. A university can become more productive while becoming less necessary to that person. No individual needs to match the strongest laboratory across all of mathematics for many such thresholds to be crossed.[42]

An alternative source of help becomes useful through gains the reader can examine. A difficult argument may be reduced, through checkable steps, to a particular unresolved lemma; a cited theorem can be checked in its original text and its hypotheses compared; a calculation can be reproduced by another method, with the assumptions varied. The remaining need for expert help becomes more definite and easier to take to another reader. Progress consists in a clearer account of what has been established and where uncertainty remains. A person can acquire greater command of an inquiry before becoming able to settle every question within it.

These conditions also shape the terms on which people learn from one another. Someone who can obtain needed explanations or criticism only from one person may have little freedom to question how that help is given. Dependable assistance elsewhere can make it easier to examine a dismissal, seek another account, or leave a discouraging exchange. The expert's abilities can remain valuable while their approval becomes less decisive. The conditions of inquiry can improve even when those who previously controlled access do not change their conduct.

This freedom matters in accepting criticism as well as questioning it. An objection may reveal that a proof fails or that a proposed connection is already known; it need not settle what remains worth learning or pursuing. Reliable assistance could help a reader reconstruct the objection and establish what it actually shows. Patient help could also make it easier to acknowledge confusion or rebuild missing background without each difficulty becoming a public assessment of ability. A claim can then be corrected or abandoned while the inquiry continues.

Autonomy can grow through relationships. Ryan and Deci distinguish it from independence: acting with a sense of volition is compatible with relying on others, and their account places relatedness alongside autonomy and competence. A teacher, colleague, or reading group can help someone acquire the means to direct their work. Dependable AI assistance can add to these relationships. The aim is room to examine and revise their terms, including the purposes the relationships help to form.[43]

A suggestion may make a new direction visible; a counterexample may change what seems worth pursuing. Self-government includes this ability to reconsider a purpose. It does not require protecting an original intention from every influence. An inquiry becomes freer when a useful challenge can change its course without placing its continued existence at the challenger's discretion.

The same assistance can make unfamiliar contributions easier for institutions to examine. Whether their practices change accordingly depends on how they use the new capacity and support the work it still requires.

III. Giving greater capacity a shared purpose

7 When the grounds for a judgment change

New capabilities matter to this argument through the work they make possible. First Proof's 2026 second batch tested four public systems on ten research problems with unpublished human solutions. Across the systems, seven problems received at least one solution judged essentially flawless or needing only minor revisions by expert referees. Failed attempts included unsupported difficult steps and defective references. The experiment established capability on specified problems; it did not assess autonomous selection of important questions or invention of research directions.[44]

Larger efforts show both the scale of the change and the resources it can require. In September 2026, OpenAI released a proposed resolution of the forced alternatives in the Navier–Stokes Millennium problem, comprising an analytical proof and a Lean formalization. Clay described the problem as apparently settled while emphasizing an unhurried evaluation. This was not a prize adjudication. The October 7 revision history of OpenAI's separate mathematical collection records

three withdrawals following a sign error in a construction used by two dependent papers, alongside repairs to other manuscripts. Internal systems, computation, and human direction contributed to the work; its resources were not thereby made available to every individual. Mathematical production outside the traditional academic career system can be consequential while creating new concentrations of capability and new requirements for examination.[45][46][47][48]

A formal proof changes part of that examination. Its certificate supplies evidence of an encoded entailment relative to specified declarations and assumptions. Whether the statement expresses the intended question, whether the result is new, and why it matters remain further judgments. Independent replay lets an inspectable artifact carry some of the burden otherwise assigned to familiarity with its author. Status cannot repair or invalidate a checked derivation, although it can still affect who interprets the statement and decides what to do with it.[49]

Libraries extend the value of the artifact. Avigad's account of the formal turn describes definitions and dependencies becoming available for reuse, refactoring, and collaboration. A contribution enters an environment in which others can locate what it assumes and build on what it establishes. That environment requires mathematical choices and maintenance; it does not emerge from a proof checker alone.[50]

The Equational Theories Project demonstrates a concrete participation arrangement. More than fifty people with varied academic and professional backgrounds investigated relations among 4,694 equational laws. Public tasks, reviewed submissions, discussion, and automated replay supported their cooperation. Lean certified generating facts; external programs computed the full graph's closure. The infrastructure was chiefly formal and automated-theorem-proving work, with limited LLM assistance. It supplies a collaborative precedent rather than a measured reduction in status bias.[51]

A contributor can find a task, produce an inspectable result, receive review, and enter a shared project. Experts remain responsible for specifications, onboarding, and maintenance. Personal acquaintance need not bear the entire burden of establishing reliability. The project creates a route through which an unfamiliar participant's work can become useful under demanding standards.

A barrier defended by the cost of examining unfamiliar work needs renewed justification when that cost falls; if the remaining cost no longer warrants it, another reason is needed. A claim may still be wrong, unoriginal, poorly explained, or irrelevant to the project. Those grounds require their own assessment.

The obligation works in both directions. Cheaper candidate arguments can make interpretation and review relatively scarcer, increasing reliance on familiar names. New requirements and additional resources may be justified when the burden rises. Reader assistance may make an initial examination easier without making complete verification cheap. Procedures should track the actual work required, rather than treating either technological enthusiasm or an inherited practice as sufficient justification.

Scarcity at the point of assessment also has a history. When a subject depends on a few people able to explain or check its arguments, new questions return to the same bottleneck. An exposition that makes a decisive reduction teachable, or a maintained library that makes its dependencies

inspectable, can enable additional readers to undertake that work. Such an investment uses resources that could serve the present queue; its effects take time. Institutions influence whether it is sustainable through what they reward and whom they support. A culture can neglect the work that would enlarge its capacity for judgment, then invoke the resulting shortage as a reason to keep the channels of judgment narrow.

There is a reciprocal test for the standards invoked. If understanding is used to resist a theorem-only conception of AI achievement, explanation and maintainable tools should receive recognition in human careers too. If authors must disclose assistance and accept checking, evaluators should take responsibility for their own judgments and tool use. If a company claims a machine achievement, its account should preserve the human mathematics, libraries, scaffolding, and revision on which the result depends. Greater capability makes the organization of responsibility more consequential.

8 Credit as an intellectual record

A checked result and an account of how it became possible answer different questions. Credit directs readers toward ideas, methods, and people whose work they may need. It also affects which kinds of labor remain viable. When recognition compresses a collective achievement into a single victorious endpoint, the loss extends beyond a distribution of rewards: part of the record from which later inquiry could learn disappears.

Thurston's account of mathematical progress connects this issue to his own experience. Early success in foliations could discourage others from participating; his later practice emphasized communicating ways of thinking. He acknowledged the unusual trust and freedom his standing afforded him. The institutional question is how to make the labor of creating shared understanding visible and sustainable beyond those who already possess that freedom.[52]

Shimura's letters to Richard Taylor, reproduced in *The Map of My Life*, identify a related loss. He complained that later summaries narrowed the scope of his earlier work and that readers repeated them without returning to the originals. A subsequent letter explained the mathematical considerations behind his modularity conjecture. His testimony concerns the reception of his work, but it points to a general use of attribution: preserving the reasons that made a conjecture plausible, rather than remembering only the endpoint through which it became famous.[53]

A useful account distinguishes a conjecture, a crucial reduction, a proof of the remaining claim, a correction, and an explanation that made a method reusable. Different roles can deserve different credit. Responsibility for the final account is also distinct from having discovered everything it contains. These distinctions let readers reconstruct dependencies and give participants reasons to share formative ideas before a finished result exists.

They remain possible when a model contributes. In the 2026 primitive-set preprint by Alexeev and collaborators, model output suggested a Markov-chain method that human constructions and checking developed into several proofs, including two previously open conjectures and the revised Banks–Martin conjecture. The authors distinguish model assistance, human contributions, discarded suggestions, and responsibility for final proofs. An unfamiliar form of collaboration calls

for a more precise record of the work, not an undifferentiated attribution of agency to whichever participant makes the announcement.[54]

Recognition can be arranged to support exchange. Daubechies recalls Bell Labs bringing colleagues from different departments together, hosting the unsolved-problem seminar through which she began working with Lagarias, and giving each collaborator more than half credit for joint papers in salary reviews. Those practices made conversation easier and reduced a reason to experience cooperation as dividing a fixed reward. They supported generosity through ordinary working conditions.[55]

Explanation and maintenance need a place in those conditions. Cohn's discussion of AI-generated mathematical technical debt describes valid but poorly explained arguments that leave substantial work for others, often with less recognition for the subsequent explanation. He also welcomes responsible newcomers documenting AI-assisted work. Sanderson emphasizes explanations that reconstruct why a construction was introduced and how someone could have arrived at a method. Both accounts direct attention to work that changes what others can understand and use.[56] [57]

Its effect can be assessed without pretending to a universal score. Can readers locate the decisive idea, identify its limits, and adapt it elsewhere? Does a library make a dependency inspectable? Has a corrected argument become usable by another project? Evidence of these effects supplies reasons for recognition. A citation count records specified events; it does not supply the missing explanation of significance. The IMU–ICIAM–IMS report on citation statistics details why journal averages and short citation windows cannot substitute for assessment of individual mathematical work.[58]

Avigad's account of young tool builders whose work receives insufficient recognition within mathematics identifies the practical stakes. Institutions seeking better reasoning technology cannot consistently treat its construction as irrelevant to mathematical progress. A field can depend on a form of work while making it difficult for its practitioners to remain in the field.[59]

Understanding itself need not be divided like a prize. Another person's learning an argument does not diminish what the first reader understands; a newcomer's question can disclose something both had missed. Attention, recognition, and resources remain limited. An appointment or grant must be allocated, but the explanations, examples, and reusable tools developed through that support can serve people who were not selected. Making these benefits usable beyond the selected group requires further work. Supporting it extends the intellectual reach of a limited allocation.

9 Making participation substantial

A route into shared work needs a defined demand on scarce attention. A project can state the questions it is prepared to examine and reserve review for bounded contributions: a reproducible example, a missing argument supplied, or a result meeting specified conditions. Clear prerequisites and examples let an unfamiliar contributor prepare work whose relevance can be assessed without reconstructing an entire private theory. A first assessment can determine what warrants deeper examination, while distinguishing a defect in the work from a shortage of review capacity. Someone

must undertake that assessment and integrate what proves useful. The arrangement gives evidence from an unfamiliar participant a way to affect the decision to invest further attention.

Polymath16 illustrates how examination can gather around inspectable work. In 2018, de Grey presented a finite graph establishing that the chromatic number of the plane is at least five. Independent checks exposed an error in a proposed smaller graph while confirming a larger construction. He corrected his code; participants shared checks and developed simplifications. The initial contribution gave people a reason to cooperate, and visible correction strengthened the grounds for continuing. Public availability became an exchange through particular people's work.[60] [61]

For ordinary evaluation, the grounds of a decision should be portable enough to examine. A missing hypothesis, failed estimate, counterexample, comparison with prior work, or clearly stated project criterion gives another competent reader something to assess. A judgment about an entire framework may require more context, but can still identify what fails to connect or which promised consequence remains unsupported. Experience can be tacit; a consequential decision needs enough explicit grounds to make correction possible.

The IMU's journal recommendations already call for impartial evaluation, grounds for decisions, and appeals alongside screening and editorial judgment. An appeal is useful when an additional competent assessment can change the outcome. Repeating a procedure without making its grounds examinable can preserve a decision without testing it. Independent readers and disclosed conflicts help, with the arrangements scaled to the consequence of the decision and the feasible work of review.[62]

No procedural label guarantees fairness. In a randomized ecology-journal trial, blinding reduced some location-related advantages but produced no detected difference in its effects between male- and female-authored papers; editors still knew identities. The result concerns that procedure and setting. A reform should be assessed through what it changes, including the barriers it leaves untouched.[63]

New collaborations also need agreed purposes. Avigad's account of sphere-packing formalization describes a project that accepted AI contributions but worried that a company's separate completion would be presented as the whole achievement. Human scaffolding, appropriate credit, and reusable infrastructure mattered. He records a repair as well: the company cooperated on attribution and promised open work on improving the code, leaving the community cautiously optimistic. Correctness, credit, and maintainability needed distinct terms, through which the participants could negotiate a useful outcome.[64]

The EMS code assigns AI disclosure and responsibility to authors and referees, with protections for confidential material.[24]

Collective support makes such standards practicable. Gowers's account of withdrawing from Elsevier explains why people can keep sustaining an arrangement they criticize: declining a request feels like letting down an editor or author they know. Public coordination can make another choice easier. Reform needs ways for people to act together without abandoning the collaborators on whom their work depends.[65]

The same need appears in libraries, review systems, explanatory resources, and shared records. They require maintenance, time, and recognized responsibility. Depending indefinitely on unpaid enthusiasm for essential infrastructure gives expanded participation a weak foundation. Access to useful computing also needs provision if the opportunities are to reach beyond those who can afford private resources.

AI firms are institutions too. Academic authority can decline while control over computation becomes concentrated elsewhere. Affordable access, inspectable arguments, and records that can be retained and continued through other tools give people more room to disagree and change direction. The means of continuing matter at this new point of dependence just as they mattered in the earlier encounter with a teacher or department.

Several forms of inspection should remain available. Fields with mature formal libraries have opportunities that fields with new definitions and little infrastructure do not yet possess. Making one technical format a universal entrance requirement could reproduce exclusion through a new convention. Appropriate evidence can include formal certificates, careful informal arguments, and explanations connecting an encoded statement to the intended mathematics.

Developing more evaluators is also an educational task. Tao's account of AI-era mathematics connects understanding with responsibility, attribution, and education. The Leiden Declaration, endorsed by the IMU, joins these commitments with accessible participation and shared infrastructure. Their practical meaning lies in the work people become able to undertake: formulating questions, challenging claims, revising arguments, and helping others enter a subject. Funding, appointments, and project standards determine how far those commitments can be exercised.[66] [67]

Entry materials develop competence; review makes claims examinable; accurate attribution preserves contributions; resources sustain the people doing this work. Each supplies something the others need. An invitation without reviewers can overwhelm a project. Checking without explanation can leave its results difficult to use. Praising maintenance without supporting it can erode the capacity on which the rest depends.

10 An inquiry with somewhere to go

A different example begins with two familiar integrals. We can compute their values and still ask which changes of coordinates preserve the information they carry together:

$$\int_1^2 \frac{dx}{x} = \log 2, \qquad \int_{|z|=1} \frac{dz}{z} = 2\pi i,$$

with the circle oriented counterclockwise. In Brown's realization framework, the relative cohomological object $H^1(\mathbb{G}_m, \{1, 2\})$ has, in suitable bases, the comparison matrix

$$\begin{pmatrix} 1 & \log 2 \\ 0 & 2\pi i \end{pmatrix}.$$

The rows correspond to the interval and circle. Integrating the differential of the coordinate gives the first column. Although the differential of the coordinate is exact on the punctured line, its primitive takes different values at the two marked endpoints, so its relative cohomology class need not vanish. The two displayed integrals give the second. The cohomological object joins two simpler pieces; a split object with the same graded pieces admits a comparison matrix with zero in the upper-right entry. In the rational mixed Hodge realization, a splitting would require the path integral to be a rational multiple of the loop integral. The first is positive and real, the second purely imaginary, so no such splitting exists. The pieces agree, but their attachment carries information of its own.[68] [69]

A concrete attempt is to replace the path class by itself plus a rational multiple of the circle class. This keeps its boundary and changes its logarithmic period only by a purely imaginary number, so the positive real part cannot disappear. The restriction to rational coefficients and compatibility with the Hodge filtration matter: a valid identification of complex vector spaces need not identify the structured objects. Assistance now has a definite job: explain those restrictions, work through the allowable changes, and locate the exact step where a tempting simplification fails.

The failed change leaves a specific fact: adding any allowed loop correction cannot remove the real part of the path integral. It also locates the assumption on which that conclusion depends. If arbitrary complex multiples were allowed, the upper-right entry could be removed; the obstruction belongs to the structured comparison, not to an unqualified matrix problem.

That suggests a comparison with the coordinate-repair example. There, polynomial coefficients specify the allowed corrections. Here, the path boundary and the realization data restrict them. In each case one can first list the admissible changes and then ask whether any of them removes the residual.

A reading group can begin at any of these steps: check an admissible change, test a different example, or explain why relaxing a condition changes the answer. The resulting notes give the next reader a definite claim to examine and a place to continue.

A reader returns to a familiar integral and sees, this time, why a comparison contains more than a list of dimensions. Another works through a proof with a group until its hypotheses become intelligible. Someone brings a failed argument and leaves with a clearer question. The explanations were already available. Now they can be followed, questioned, and carried into another conversation.

Notes and sources

1. Gowers, *The Two Cultures of Mathematics*, pp. 1–2.

2. Poincaré, *Intuition and Logic in Mathematics*, §V.

3. Poincaré, *Mathematical Creation*, in *The Foundations of Science*, p. 385.

4. De Regt, *From Explanation to Understanding*, §3. The source concerns scientific understanding and its standards; the mathematical capacities described in the essay are a considered application, not an empirical finding from that paper.

5. Tao, *What Is Good Mathematics?*.

6. Silver and collaborators, *Mastering the Game of Go Without Human Knowledge*.

7. Shin and collaborators, *Superhuman Artificial Intelligence Can Improve Human Decision-Making by Increasing Novelty*. Observational analysis of more than 5.8 million professional Go moves. Novelty is departure from recorded move sequences; decision quality is assessed by the study's AI measure.

8. Sadler and Regan, *Game Changer*, preface.

9. Teuber, interview on AI and Go study. Teuber's account as an amateur 6-dan, teacher, and AI Sensei cofounder; an explanation of teaching practice, not a controlled learning experiment.

10. Rockmore, *Should the Classroom Be More Like the Gym?*.

11. Grothendieck, *Récoltes et Semailles*, introductory letter §6. Introductory letter §§6–7, on competitive mathematics and formative work; "Motives—or the heart within the heart," on cohomological realizations.

12. Howard Straubing, "A combinatorial proof of the Cayley-Hamilton theorem", *Discrete Mathematics* 43(2–3) (1983), 273–279. The original combinatorial proof.

13. Doron Zeilberger, "A combinatorial approach to matrix algebra", *Discrete Mathematics* 56 (1985), 61–72, §3, pp. 66–67. This presents Straubing's walk-and-cycle proof and its reversible cancellation.

14. Bryan and collaborators, *The Motivic Class of the Space of Genus 0 Maps to the Flag Variety*, §1.2 and Appendix A. Version 2, 11 September 2026. The authors describe iterative human scaffolding and specialized internal systems, human recognition of a general strategy, and human responsibility for the final text. The original scaffold was not solved unaided by the systems.

15. Bastani and collaborators, *Generative AI Without Guardrails Can Harm Learning*. Randomized high-school mathematics study. The 17 percent loss is on an unassisted assessment after the unrestricted interface; the guarded tutor largely avoided the loss without establishing an unassisted-exam gain.

16. Kestin and collaborators, *AI Tutoring Outperforms In-Class Active Learning*. Randomized crossover design for two physics lessons, with expert-written solutions and structured tutoring. Immediate lesson outcomes do not establish an effect over an entire course of study.

17. F. William Lawvere, "Explicit Foundational Concepts in the Teaching of Mathematics", English scan, pp. 621–626, especially pp. 621–622. The family archive attributes the English publication to Fabio Minazzi's edited volume (Rome, 2007); the 2001 date concerns earlier Italian conference context.

18. Doron Zeilberger, Opinion 29, 24 September 1998. Autobiographical account of the Bass interruption, differing hypotheses, and subsequent work with Dominique Foata.

19. Évariste Galois, unfinished preface written at Sainte-Pélagie, December 1831, English text at MacTutor. French text in Tannery's 1908 edition, pp. 25–27. His polemical account is attributed, not treated as an impartial finding.

20. Craven, *Where Will All the Papers Go?*.

21. JMLR, editorial-board statement.

22. Gowers, *Discrete Analysis Launched*.

23. Robert K. Merton, "The Normative Structure of Science" (1942), reprinted in *The Sociology of Science*, edited by Norman W. Storer (1973), pp. 267–278, especially pp. 270–278.

24. European Mathematical Society, *Code of Practice for Mathematical Publication*, approved 19 March 2025. Referee duties 3–6; author duties 3 and 8; editor duties 1 and 3. A professional standard, rather than a measurement of compliance.

25. Jürgen Huber et al., "Nobel and novice: Author prominence affects peer review", *PNAS* 119 (2022): e2205779119. Primary full text. The anonymous-invitation comparison yielded 102 reports with the prominent name, 110 anonymous, and 101 with the early-career name; outcomes are recommendations.

26. Simcoe and Waguespack, *Status, Quality, and Attention*. Natural-experiment evidence about attention and publication at the IETF; names remained in the full drafts. The contrast with working-group drafts does not isolate scrutiny as its sole cause.

27. Robert K. Merton, "The Matthew Effect in Science", *Science* 159 (1968): 56–63, doi:10.1126/science.159.3810.56. Separate discussions concern credit, communication, and resources.

28. Thijs Bol, Mathijs de Vaan, and Arnout van de Rijt, "The Matthew effect in science funding", *PNAS* 115 (2018): 4887–4890. The subsequent-funding comparison excludes the original award; see the threshold design and analysis of later applications.

29. Caroline Series, interview in *EMS Magazine* (2022). Her comparison of Berkeley, Cambridge in 1977–78, and Warwick includes subsequent changes in the profession.

30. "Accidental activist: An interview with Lenore Blum, Part 1", *Berkeley Science Review*, 21 May 2020. The entry encounter, Singer's intervention, and later faculty treatment are distinct parts of her recollection.

31. G. H. Hardy, *A Mathematician's Apology* (1940), §§7 and 29. Mixed motives and legitimate ambition are distinguished from misconduct.

32. J. Peter May, "Memorial Address for J. Frank Adams", *The Mathematical Intelligencer* 12(1) (1990): 40–44; delivered 29 April 1989. Competitiveness, p. 40; the intimidated student and attempts to help, pp. 41–42; care for exposition, p. 41.

33. Alexander Grothendieck, *Récoltes et Semailles*, Part I, §40, "La mathématique sportive" (mid-1980s). Mateo Carmona's provisional transcription; §40 begins at PDF p. 256. Used for self-examination, without adopting separate allegations.

34. Mike Shulman, comment of 15 October 2010, in John Baez's *Vanity and Ambition in Mathematics* (14 October). Shulman's criticism includes himself.

35. Vladimir Voevodsky, "The Origins and Motivations of Univalent Foundations", Institute for Advanced Study (2014). First-person account of errors, delayed correction, reputation, and the development of his verification program.

36. Grothendieck to Brown, 14 June 1983, p. 47. The letter discusses inhibition of speculative thought and a planned exploratory book on motives. It is Grothendieck's account of the intellectual climate.

37. Grothendieck's Crafoord Prize refusal, published in *Le Monde*, 4 May 1988. Grothendieck's diagnosis and reasons for declining the prize are attributed to him; they are not treated as measured prevalence of misconduct.

38. Rachel Crowell, "Conducting the Mathematical Orchestra From the Middle", interview with Emily Riehl, *Quanta*, 2 September 2020. Final question on social issues and questions on Women in Topology and discrimination. Her positive assessment is explicitly experiential.

39. Illusie and collaborators, *Reminiscences of Grothendieck and His School*, p. 1115.

40. André Joyal, "Democracy and Mathematics", letter dated 1 October 2010, guest post published 4 October by John Baez. Final mathematical paragraphs distinguish specialization from artificial barriers and propose collective accessible knowledge.

41. Mike Shulman, comment of 26 May 2015, 6:25 pm, on *Voevodsky trending*. Appreciation of leadership accompanies his rejection of a portrayal of the field as one person's program.

42. Coase, *The Nature of the Firm*, on relative costs of organization. A comparison with the relative costs of organizing activity inside and outside an institution. Coase did not study AI-assisted mathematics; the local-sufficiency argument is the essay's application.

43. Ryan and Deci, *Self-Determination Theory and the Facilitation of Intrinsic Motivation*, pp. 70, 74.

44. Mohammed Abouzaid, Nikhil Srivastava, Rachel Ward, and Lauren Williams, *First Proof Second Batch*, report dated 10 June 2026, submitted 16 June; §§1.1–1.3 and evaluation sections. Each submitted solution received two or three independent expert assessments. The seven-of-ten figure is across systems.

45. OpenAI, *On the Navier–Stokes Millennium Prize Problem*. 8 September 2026. The announcement concerns the forced alternatives C and D; it supplies an analytical proof and formalization. This essay reports the proposed result and does not constitute an independent proof audit.

46. Clay Mathematics Institute, September 11 announcement. 11 September 2026. The announcement describes apparent settlement while retaining the Institute's deliberately unhurried evaluation process; it is not a prize decision.

47. OpenAI, October mathematical collection. Repository checked 12 October 2026: 719 manuscripts in 372 result families, with different verification states. A collection count is not a count of independently certified results.

48. OpenAI, collection revision history. The 7 October entry records three withdrawals after a sign error in a shared construction, and repairs to other proofs. Dependence among manuscripts matters to interpreting a large collection.

49. Lean language reference, "Validating a Lean Proof", checked 11 October 2026. Validation concerns the encoded statement and its assumptions; independent replay checks the supplied proof artifact.

50. Jeremy Avigad, "Mathematics and the formal turn", *Bulletin of the AMS* 61 (2024): 225–240; preprint revised 4 November 2023. §§4–10 and 17 discuss libraries, dependencies, refactoring, collaboration, and machine learning.

51. Matthew Bolan, Joachim Breitner, Jose Brox, et al., *The Equational Theories Project: Advancing Collaborative Mathematical Research at Scale*, version 2, 16 December 2025. §§1.2–1.3, 3–4, and 11 distinguish the result, certificates, workflow, and limited LLM assistance. A collaborative precedent, rather than a measured reduction in status bias.

52. William P. Thurston, "On Proof and Progress in Mathematics", *Bulletin of the AMS* 30 (1994): 161–177. Especially §§4–6, on proof, incentives, community, and his own contrasting experiences.

53. Goro Shimura, *The Map of My Life* (Springer, 2008), Appendix A3, pp. 176–180: letters to Richard Taylor, 25 November and 12 December 1994. Publisher's PDF. His account of the reception of his work.

54. Boris Alexeev et al., *Primitive sets and von Mangoldt chains: Erdős Problem #1196 and beyond*, preprint, 1 May 2026. Introduction, §1.1, and §11 distinguish mathematical results and human/model contributions. One author is affiliated with OpenAI; the work includes a new proof of already solved Problem #164.

55. Allyn Jackson, interview with Ingrid Daubechies, *Celebratio Mathematica*, conducted in late 2018, section "Babies and Bell Labs." The organizational practices are Daubechies's recollections.

56. Henry Cohn, "The technical debt of AI-generated mathematics", 15 September 2026. Guest essay hosted by Terence Tao, on explanation, responsibility, and the work left to other researchers.

57. Grant Sanderson, "If math is more than proof, we need to better celebrate the rest of it", 18 September 2026. Guest essay hosted by Terence Tao; an argument about values and recognition.

58. Robert Adler, John Ewing, and Peter Taylor, *Citation Statistics*, joint IMU–ICIAM–IMS committee report (2008), especially pp. 6–7 and 10–12. Journal citation measures and individual assessment are distinguished.

59. Jeremy Avigad, *What is mathematics now, and what should it be?*, 24 August 2026, especially §§3–5. The discussion of recognition for tool builders is an attributed professional diagnosis, rather than a population study of hiring.

60. Polymath16, discussion and corrections.

61. Mixon, independent checking.

62. IMU Committee on Electronic Information and Communication, *Best Current Practices for Mathematics Journals*, June 2021, §§1.1–1.3. Recommendations on screening, evaluation, decision grounds, conflicts, and appeals.

63. Charles W. Fox, Jennifer Meyer, and Emilie Aimé, "Double-blind Peer Review Affects Reviewer Ratings and Editor Decisions at an Ecology Journal", *Functional Ecology* 37(5) (2023): 1144–1157. Three-year randomized trial: 3,739 submissions, 1,432 reviewed. §§2–4 distinguish ratings, revision decisions, imperfect blinding, and inferred gender.

64. Jeremy Avigad, *Mathematicians in the Age of AI*, dated 3 March 2026, submitted 4 March, §2, especially its concluding paragraph. His sphere-packing account records both disagreements over purposes and subsequent cooperation on credit and open revision.

65. Timothy Gowers, "Elsevier — my part in its downfall", 21 January 2012. His account of refereeing obligations explains a coordination difficulty and the value of public collective action.

66. Terence Tao, *Mathematics in the age of AI*, submitted 17 August 2026, based on his 24 July ICM lecture. Especially the discussion of mathematical purposes, proxy measures, and §8 on responsibility, attribution, and education.

67. *Leiden Declaration on Artificial Intelligence and Mathematics*, 2 June 2026, doi:10.5281/zenodo.20302944. See commitments on understanding, verification, participation, attribution, and public infrastructure; endorsement establishes a stated institutional commitment.

68. Brown, *Notes on Motivic Periods*, §5.3. The logarithm example uses relative cohomology and its comparison data. The nonsplitting argument in the text concerns the rational mixed Hodge realization, and invokes no unproved period conjecture.

69. Burgos Gil and Fresán, *Multiple Zeta Values: From Numbers to Motives*, Theorem 2.154, pp. 111–112.

73. Falk Eisermann, "A Golden Age? Monastic Printing Houses in the Fifteenth Century", in Benito Rial Costas (ed.), *Print Culture and Peripheries in Early Modern Europe* (Brill, 2013), pp. 37–67, especially pp. 65–67. Noel L. Brann, "A Monastic Dilemma Posed by the Invention of Printing", *Visible Language* 13.2 (1979), pp. 150–167, especially pp. 151–157, on Trithemius's copying argument, mixed library, 1494 edition, and advocacy of printing.