

Will greed save mathematics (mathematicians)?

Artificial intelligence, the costs of research, and inequalities in access to knowledge

«Greed, for lack of a better word, is good, greed is right, greed works.»

Gordon Gekko, *Wall Street* (Oliver Stone, 1987).

1. From solving a problem to the end of a profession

On September 8, 2026, OpenAI announced a proof of finite-time singularity formation for the three-dimensional Navier-Stokes equations with a smooth external force, accompanied by a formalization in Lean [1]. Three days later, the Clay Mathematics Institute referred to an apparent solution, naturally leaving both confirmation of the result and the question of credit to subsequent verification [2]. The result concerns the case with a forcing term (as included in the Millennium Problem); global regularity, or its failure, in the absence of such a term therefore remains an open problem [3].

In recent days, with impressive speed, the announcement of an extraordinary result has turned into a funeral for the profession. Mathematics is over, mathematicians will no longer be needed, soon a machine will do everything. We may indeed be entering a new phase, far more radical than the previous ones; it would be foolish to deny that possibility. But the logical leap from «a machine has solved this problem» to «mathematicians will no longer be needed» seems enormous to me. Even granting that the proof is fully confirmed and that similar systems soon solve other major problems, at least three questions remain: how much does it cost to make them work, who can afford it, and for what purposes will those who control them have an interest in using them?

An exceptional campaign built around a famous problem tells us, for now, a great deal about technical capability and much less about everyday research. We do not know whether it would make sense to organize thousands of projects in the same way, nor whether the most efficient way to do mathematics will be to leave everything to the machine or to use it to multiply the work of competent people. These are economic and organizational questions, not secondary details. They are often overlooked in the most catastrophic discussions of these days.

Moreover, automating some tasks does not automatically mean that a profession disappears: economists have long emphasized that what matters is both the activities that are eliminated and those that are made possible or more productive [4]. Of course, this is no insurance policy for us. But mathematics is not a finite list of problems with boxes to tick. A major theorem closes one question and almost always opens others; one must choose which questions are worth asking, invent concepts, recognize connections, simplify, explain, teach. Perhaps a growing part of this work will be assisted or even automated. At least today, however, I see no serious reason to regard the entire scientific process as already automated.

2. The cost of a theorem

OpenAI's account mentions about 10,000 simultaneous agents in the decisive group, 88 hours of research, and another 17 for formalization and verification. For Navier-Stokes, it reports about 130 billion output tokens, out of about 300 billion for the full set of problems addressed

[1]. These are figures worth keeping in view. Whatever one thinks of the result, we are not talking about a long session with the chatbot each of us uses on our own computer.

A token is a unit through which the system represents and processes content; it does not necessarily correspond to a word. APIs make it possible to use models from programs, including by running many instances together, and pricing schedules distinguish among several components of usage [5]. The price of a personal ChatGPT or Claude subscription therefore says very little about the cost of a genuinely autonomous research campaign. The interface may look the same, but underneath it the scale of computation can differ by several orders of magnitude.

On the figure of 80 million dollars, I prefer to be cautious: I use it here as a cost scenario, not as an established accounting of the Navier-Stokes operation. Even if we had a certain figure, we would still need to understand what it included. Model training, inference on the problem, failed attempts, human labor, infrastructure, and the final price charged to a customer are not the same thing. It would be wrong to assign to a single theorem the entire cost of a model reused in a thousand other ways; it would be equally wrong to count energy and infrastructure twice.

A very rough calculation at least helps clarify the point. At the date consulted, the public standard price of GPT-6 Astra for short-context output is 50 dollars per million tokens [5]. Applied to 130 billion tokens, this would amount to 6.5 million dollars for output alone. This is not an estimate of the cost of Navier-Stokes: the model used was internal, other cost items are missing, and the commercial price is not the same as the industrial cost. The calculation does not confirm the 80 million figure; it merely shows how easy it is to speak of the «cost of the theorem» without actually having the experiment's accounts.

For an organization that funds research, the interesting quantity is different: on average, how much does it cost to obtain results that are verified and reusable? One would also have to include failed attempts, human labor, external tools, and preservation of the materials. The spectacular success we see is inevitably selected. To understand whether this mode of research is genuinely generalizable, we would also need data on the paths that led nowhere.

None of this diminishes the importance of the result in the slightest. If a machine solves a problem that has resisted for decades, we may well decide that an enormous expense is worth it. But saying that it is worth doing once is different from saying that this will become the normal way of working on every lemma, every PhD thesis, every still-confused intuition. It is precisely in this distance between exception and routine that, in my view, part of the future of the profession will be decided.

3. The value of mathematics to a technology company

Saying that OpenAI and Anthropic «do not care about mathematics» is deliberately blunt and, taken literally, unfair. I do not know what personally motivates the researchers who work there, and I have no reason to doubt that many of them are sincerely interested in scientific problems. But a company remains a company. The intrinsic importance of a theorem is not enough, by itself, to determine where billions of dollars are spent. No board of directors has a duty to support all of pure mathematics for decades, including areas that make no headlines and promise no economic return.

OpenAI explicitly writes that publishing the result also serves to document the progress of its models [1]. It is therefore difficult not to see it also as a public demonstration of capability. Solving a famous problem counts as research and, at the same time, as advertising: it tells customers and investors that the technology can do something that until the day before seemed out of reach. The two things can coexist perfectly well.

The economic return may come from markets that have very little to do with mathematics. An expenditure that could never be recovered by selling theorems to universities may become reasonable if it persuades banks, industries, governments, or large companies to buy services. But it is precisely here that I see a limit to the darkest predictions. The first Millennium Problem solved by a machine has enormous symbolic value; the ten-thousandth technical problem in a highly specialized subfield much less so. I see no reason why the investment should remain of the same order.

There is also a simple fact: once a proof is published, it largely becomes a good that others can use without paying its discoverer each time. Uncertainty and the difficulty of appropriating the economic value of research are classic themes [6]. For a company, it may be far more natural to sell access to the machine that searches for the next theorem than to sell the theorem already found. We might therefore have public knowledge and, at the same time, highly concentrated means of discovery.

Paradoxically, this too seems to me a reason not to indulge too much in catastrophism. The market has no interest in spending unlimited resources to exhaust mathematics, even assuming that it were technically possible to do so. If a problem produces no return, the computation will be used elsewhere. That problem will remain there, and someone will still need a scientific reason to work on it. Universities, public institutions, and research communities do not become less necessary; if anything, it becomes more important to understand what we want to fund independently of its immediate commercial value.

4. The price of tokens and the actual cost of research

Should we therefore expect the price of tokens to rise continuously? I am not at all sure. The price of already available performance can fall rapidly: Stanford's AI Index reports, between November 2022 and October 2024, a reduction by more than a factor of 280 in inference cost for performance comparable to GPT-3.5 on the MMLU benchmark [7]. This does not measure the cost of mathematical research, but it is enough to make any linear prediction of the form «tomorrow every token will cost more» suspect.

The problem, however, is that what matters to us is the total expenditure required to obtain a result. If an individual token costs half as much but ten times as many are needed, we pay five times more. Parallel agents, long contexts, repeated verification, and alternative strategies can quickly consume efficiency gains. The opposite can also happen: a better model might find the right path sooner. This is precisely why, in my view, it makes little sense to fixate on a single line in the price list.

One indication comes from Anthropic itself. In its 2025 account of its research system, the company reported that multi-agent architectures used about fifteen times as many tokens as ordinary conversations [8]. This is a figure for that system and those tasks, not a universal law. But it makes very concrete the distance between the assistant we use every day and a procedure that autonomously explores many paths for hours or days.

We might also pay more without seeing a dramatic increase in the nominal price. The best model could be moved to a different plan, usage limits could be tightened, priority access or long contexts could be sold separately. Price lists already distinguish among models, token types, and modes of use [5, 9]. For someone planning a department's spending, the price of one million tokens is therefore only a small part of the problem.

Companies must recover investments and generate margins. This makes price increases or restrictions plausible when demand is high and alternatives are few; it does not make them inevitable, because competition and technical progress can push in the opposite direction. I do not know which of the two forces will prevail. I do know, however, that scientific dependence on commercial conditions decided elsewhere is a real problem, especially when changing provider means discarding months of procedures, integrations, and working habits.

And then there is the physical side, which the word «cloud» tends to make us forget. In its 2025 Energy and AI report, the International Energy Agency projected global electricity demand from data centers of about 945 terawatt-hours in 2030, more than double the 2024 level [10]. This is a scenario, not a certainty, and it concerns far more than mathematics. But it reminds us of something obvious and important: behind an apparently immaterial answer there are power plants, grids, chips, cooling, and maintenance. Abundance on the screen does not eliminate scarcity in infrastructure.

Against this background, I find today's situation almost surprising. For a mathematician who knows how to use these tools well, a personal subscription on the order of one hundred euros per month can be an exceptional bargain, almost a «gift» from OpenAI and Anthropic to the community. Both companies offer an individual tier at 100 dollars per month; I therefore use one hundred euros as an order of magnitude, not as an identical price in every country [21, 22]. The gift, of course, is not philanthropy: it is the possible disproportion between what we pay and what we are able to obtain from it scientifically. And for a precarious researcher or a student, it should be said, even that amount may already be significant.

I would not, however, assume that this balance will last. By paying a subscription we buy access, with limits and conditions that can change; we do not own the machine. Anthropic, for example, makes clear that plans and prices may change [21, 22]. Capabilities that today are included in a personal subscription might tomorrow require a much more expensive plan or usage-based payment. Perhaps prices will instead continue to fall. We do not know. That uncertainty itself seems to me the point.

There is also a more radical possibility: personal access to the best systems might simply be withdrawn. The retirement of individual models is already part of the normal life cycle of services; Anthropic explicitly notes the consequences for ongoing research and comparative studies [23]. Replacing an old model does not, of course, mean closing the door on research. But it reminds us that what we can use today is not necessarily a tool we will be able to keep tomorrow. The most capable system could remain internal or be reserved for a few institutional clients.

Chess and Go, to which I return below, at least make this scenario imaginable. Deep Blue was retired; AlphaGo stopped competing after demonstrating what it had been designed to demonstrate [12, 13]. In those cases, access to game-playing programs certainly did not disappear, and Stockfish is there to remind us of that [17]. For mathematics the point would be different: we might lose access to a genuinely useful system without having an equivalent re-

placement. At that point the question would inevitably become political: who guarantees access, at what cost, and according to what criteria?

For this reason, with all the necessary qualifications, I regard the present moment as a kind of golden window. We can try tools of a power that was unthinkable until very recently without first having to win a major grant. I do not know whether this will last two years, ten, or much longer. The most sensible reaction, it seems to me, is neither panic nor passive waiting: it is to use them now, learn what they can and cannot do, and turn temporary access into something that remains - ideas, checked proofs, understood methods, reusable code, better-trained young people. The tokens we consume disappear; what we have genuinely understood does not.

5. A new inequality among mathematicians

Mathematics has never been an egalitarian activity in terms of resources. What matters includes time free from other duties, job stability, colleagues, travel, libraries, and the possibility of being in a strong scientific environment. If, on top of all this, we add expensive computational campaigns, a new and very direct inequality appears: having an idea might not be enough; one might also have to be able to afford to test it.

The advantage could become cumulative. More funding means more attempts; more attempts can produce more results; more results can attract further funding. It is not written that things must go this way, and it would be a mistake to present it as destiny. But if we continue to evaluate only final outcomes, pretending that everyone had the same tools, we risk mistaking part of a group's wealth for individual scientific quality.

The case that concerns me most is that of a young person with a good idea and few resources, watching a richer group explore a hundred variants of it while they can check only two. In that world even the choice of problems becomes, at least in part, a budget choice. It is a serious risk. But here too I would like to avoid another automatic prophecy: cheap and widely available tools can also have the opposite effect, giving a small group opportunities that were once reserved for major centers. Much will depend on where the genuinely accessible tier of the technology settles.

We also need to separate the capacity to discover from the possibility of reading and verifying what has been discovered. A published proof can be checked without repeating the entire campaign that found it; a formal certificate may cost far less than the search for that certificate. We might therefore have widely accessible results and highly concentrated means of discovery. For this reason, I would speak of concentration of the means of research rather than a monopoly on truth.

Part of the response could be surprisingly traditional: shared infrastructure, open libraries, shared intermediate results, computing time allocated on the basis of ideas. Open models would be useful, but they are not enough: downloading software does not pay for hardware or energy. On the other hand, if a substantial share of capability remained inexpensive, we might also see a democratization of research. We do not yet know which of the two effects - concentration or diffusion - will be stronger, and it seems premature to speak as if the outcome were already settled.

6. From funding researchers to funding tokens

The risk I see as most immediate is not the cinematic one of an empty department, with machines proving theorems on their own. It is something much more prosaic: hiring one fewer young researcher in order to buy more computation. In a PRIN or an ERC grant, we might soon find ourselves allocating a substantial part of the budget to API credits, giving up a PhD student, a postdoc, or a long research visit. I am not saying that this transformation is already taking place everywhere; I am saying that it is plausible enough to deserve attention before it becomes normal.

In that case, the path of public money would also change. Instead of building expertise within the department, a growing share would go to external providers. As long as we pay, the capability is there; when the credits run out, what remains is what we have produced, but not the ability to continue. If in the meantime we have not trained people, no subscription renewal can recreate them in a few months. The issue, then, is not only how much we spend: it is how much scientific autonomy remains within institutions.

That said, here too I would not turn a risk into a destiny. If the tools greatly increase productivity, it might become rational to fund more people and more computation together, rather than necessarily choosing between them. A young researcher is not merely labor for the current project: in ten years they may pose problems we cannot imagine today, train students, recognize errors, and hold a group together. Precisely for this reason, we should avoid automatically translating the first efficiency gain into staff cuts.

The 80 million hypothesis is useful for giving an order of magnitude. If, merely for comparison, we assume a total annual cost of 160,000 dollars per researcher, we are talking about 500 person-years: fifty positions for ten years. This is not an estimate of academic salaries and does not mean that fifty people would have solved Navier-Stokes. It simply reminds us that a single very expensive campaign competes, in budgetary terms, with scientific programs lasting a generation.

The comparison, of course, cannot be settled by counting papers. A single major theorem may be worth more than hundreds of ordinary results; an automated campaign may leave behind reusable methods and software. But a research group also leaves behind trained people, judgment, scientific relationships, and continuity. The question I would ask is very simple: when the funding ends, what remains? It is a question that applies both to a machine and to a department.

Buying private services with public funds is not, in itself, a scandal. It may be the best choice. But I would want funders to have a clear view of costs, the portability of materials, the availability of proofs and code, the possibility of changing provider, and above all the effect on training. If a machine saves us half the time, the first idea need not be to halve the number of mathematicians. It could also be to double the amount of good mathematics we are able to do.

7. The accelerators of mathematics

High-energy physicists have lived for decades with a distinction that could become familiar to us: a machine may be scientifically desirable and, at the same time, too expensive to build or use as much as one would like. In 2025 CERN estimated the cost of the electron-positron phase of the Future Circular Collider at 15 billion Swiss francs, including tunnels, infrastruc-

ture, and detectors [11]. This is the price of a concrete project, not, of course, «the cost of understanding physics».

I do not want to push the analogy further than necessary. We do not know whether a larger accelerator would be enough to clarify string theory or other fundamental questions; every large machine also carries scientific uncertainty. The point is different: the fact that a technology may allow us to learn a great deal does not eliminate the question of how much we are willing to spend and which alternatives we give up funding.

The chat I have open on my computer today does not resemble a particle accelerator. A system coordinating tens of thousands of agents for months on large mathematical programs, economically, might begin to resemble one. An individual department would not need to build the «machine»: it might simply be unable to afford enough time on it. The barrier would be a price or a computation quota, not a tunnel under Geneva.

This opens up a curious scenario and, in my view, one less apocalyptic than it may seem. We might have machines capable in principle of attacking a great many problems, but not enough computation to use them on everything. In that case scarcity would return through the back door in the form of choices: which problems deserve the machine? Which can be tackled with people and more modest tools? The theoretical availability of a capability never coincides perfectly with its actual use.

And large infrastructures also have a positive side: they force us to build communities, consortia, access criteria, and shared expertise. I see no reason why a «mathematics machine» should necessarily drive mathematicians out. It might instead require people capable of formulating good research programs, choosing sensible attempts, recognizing important results, and turning them into new theory. It would be, in part, a different profession, certainly; not necessarily a nonexistent one.

8. Chess: a precedent and its limits

Chess inevitably comes to mind. Machines surpassed the champions, and twenty years later champions still exist. It is a reassuring precedent, but it should not be used as a lazy metaphor. In chess we must distinguish at least three things: the interest of a company that wants to demonstrate the superiority of its machine, the economic value of sport between human beings, and the practical possibility of banning assistance during a game. Only part of this has an equivalent in mathematics.

In 1997 Deep Blue defeated Garry Kasparov. IBM then retired the machine and, in its own account, emphasizes the reuse of the expertise developed in other areas of computing and economic activity [12]. Deep Blue did not begin a career as world champion. The demonstration had been made: for IBM, continuing to beat human beings at chess was evidently not the core of its business.

For Google DeepMind the cleaner parallel concerns Go. In May 2017, after the matches with Ke Jie, DeepMind announced that AlphaGo would no longer play competitive matches and that the group would move on to more general algorithms and other scientific challenges, while continuing to produce materials for the Go community [13]. In chess, by contrast, AlphaZero was compared with Stockfish rather than with a human champion: by then the point was to demonstrate the generality of the method [14].

This seems to me a useful lesson. Once the spectacular demonstration of superiority has been achieved, industrial interest in repeating it may evaporate. This does not mean that IBM or Google «hate» chess; it means that they have other goals. Something similar could happen in mathematics. Solving Navier-Stokes has enormous symbolic value; funding forever thousands of specialist problems that nobody outside a small community knows about is a different matter.

Naturally the boundary is not so sharp: mathematics also has a playful dimension, and chess has an obvious mathematical aspect. But there is a decisive difference: chess is a sport. The public pays to watch Carlsen play another great player, not to know the theoretically best move in the abstract. Pressure, error, time, rivalry, and the history of the protagonists matter. Google sponsored the 2024 World Championship [15]. This shows that a company may no longer have any need to prove that machines are superior and, at the same time, may find the spectacle of humans playing one another economically attractive.

Mathematics does not automatically have this safety net. A correct proof remains correct even if it was found by an artificial system; it would make no sense to reject it because it does not comply with a sporting rulebook. We may prefer a human proof if it is clearer, shorter, or more fruitful, but for mathematical reasons, not because we need to preserve a competitive category. This makes our problem more difficult than the problem in chess.

In chess, moreover, the separation between human performance and computer assistance is at least practicable. A game has a defined place and duration; FIDE rules prohibit external sources, advice, and unauthorized devices, and allow checks by the arbiter [16]. These systems are not perfect, but they make a competition between humans without an engine during play credible.

For research I see no equally natural boundary. A paper takes shape over months or years, among readings, seminars, emails, calculations, other people's ideas, and failed attempts. We can supervise an examination without computers; we cannot realistically certify that no mental step in a piece of work ever received artificial assistance. And, frankly, I am not even sure why we should want to. It makes sense to require transparency about tools and responsibility for the result; turning mathematics into a «human only» competition would seem to me a step backward.

There is also an almost paradoxical economic difference: Stockfish is free, open source, and runs on ordinary machines [17]. Human chess does not survive because the computer is too expensive; it survives because we have decided that human performance has an independent value. We cannot rely on the same convention in research. But the precedent remains useful at least for dismantling one simplification: the fact that a machine is much stronger does not, by itself, imply that human activity loses all meaning or all market value.

In short, chess makes me moderately optimistic and equally moderately cautious. It reminds us that technical superiority does not decide everything; it also reminds us that every activity survives for reasons of its own. For mathematics, the most plausible path does not seem to me to be a sanctuary without computers, but an increasingly intensive use of machines within a human community that continues to decide what is worth understanding.

9. A technological and mathematical Moon race?

There remains a question that may seem like science fiction today and perhaps much less so in a few years: will governments leave the most powerful machines in the hands of a few private companies, or will they end up, in some form, taking over at least part of their management? If the capacity to produce new knowledge became a strategic asset, the temptation of a technological - and perhaps also mathematical - Moon race would be strong. The discussion would no longer concern only the price of a service, but sovereignty: who can use the machine, on which problems, and under what constraints.

Some institutional signals already exist, without any need to speak of nationalization. In 2025 the European Commission's AI Continent Action Plan linked artificial-intelligence infrastructure to competitiveness and technological sovereignty, envisaging AI Factories and AI Gigafactories [18]. On November 24 of the same year, the U.S. Department of Energy announced the Genesis Mission, explicitly invoking the Apollo program and bringing together national laboratories, companies, and universities [19]. These are not programs for pure mathematics, but the direction is clear: advanced computing is already a matter of industrial and science policy.

Public intervention would not necessarily require expropriating OpenAI or Anthropic. A state could build its own infrastructure, fund consortia, purchase reserved shares of computation, or impose conditions when buying services. Ownership, funding, and control are different things. For a researcher, however, the decisive difference would be simple: passively accepting a provider's conditions, or having institutions capable of negotiating the purposes and rules of access as well.

A major mathematical problem could even become a symbol of national primacy. A government might spend enormous sums not because it hoped to sell the proof, but because it wanted to show that its infrastructure was superior. This is where the analogy with the Moon race seems interesting to me: an investment that would be irrational for a department may be perfectly rational for a state buying prestige, technology, expertise, and strategic independence.

This scenario naturally weakens my thesis about greed. If the state pays, the absence of a direct commercial return may not matter. Automated research could receive enormous resources. But the state, too, does not have infinite resources: it will have to choose what to fund, and not every open problem will become a national project. Here again, therefore, the extremely powerful machine does not eliminate scarcity; it shifts it to the level of political decisions.

Nor is it certain that public intervention would reduce the power of companies. It could also turn them into preferred suppliers to very wealthy public clients. The researcher might pay less out of pocket, while total spending and providers' revenues increased. Knowing who pays for the tokens is not enough: what matters is who controls the infrastructure, who sets the priorities, and how much freedom remains for those doing research.

For us the question would be very concrete. If tomorrow a personal subscription no longer provided access to the useful system, who would allow us to continue? Public infrastructure or a national agreement could be the answer. But at what cost? With how many hours per person? Under what selection procedures? It is one thing to have a free machine for ten «strate-

gic» projects; it is another to be able to ask it something on a Tuesday afternoon because we have had a perhaps foolish idea. Research also lives on this small-scale freedom.

The geography of inequality could thus shift from departments to countries. Some states might be able to afford their own infrastructure, while others become permanent customers. We might gain power and lose openness, especially if strategic priorities or confidentiality constraints came into play. Public funding, by itself, does not guarantee open science at all.

But there is also a decidedly more interesting scenario here, one that I hear invoked much less often in these days' discussion. The same public resources could build genuinely shared scientific infrastructure: transparent access, international collaboration, reusable results, and room for problems that the market would never fund. If organized well, an expensive machine could do what other large infrastructures have done: increase the number of questions that can be addressed rather than reduce the community asking them.

Naturally, machines and people would need to be funded together. And this brings me back to the issue that, paradoxically, concerns me more than the extinction of the profession: how will we recognize the young people in whom to invest, if the results with which they present themselves increasingly depend also on the quality of the tools they had available?

10. The real problem: recognizing and recruiting young mathematicians

If I had to identify a problem that I genuinely find difficult, it would not be deciding whether mathematicians will still exist in twenty years. It would be understanding how we will select young ones. Evaluating a paper and evaluating the person who presents it are not the same thing. We may become increasingly good at checking whether a result is correct; at the same time, it may become harder to understand how much that result tells us about the candidate's scientific independence.

In fact, it has never been simple. Every paper contains the influence of collaborators, mentors, available time, the group in which one works, even the luck of encountering the right problem. Artificial intelligence, however, adds a new and potentially very powerful factor. Someone with twenty years of career behind them can also be judged on their history; a twenty-five-year-old has to convince us with very few signals, precisely while those signals risk becoming more ambiguous.

Imagine two candidates. One works in a wealthy group, with access to a system that explores hundreds of strategies; the other has modest resources. The first presents a stronger result. Is that person the better mathematician? Perhaps yes, perhaps no. It would be absurd to penalize them for having known how to use advanced tools well; it would be equally absurd to pretend that the tools do not matter. That is precisely the problem: from publications we may no longer be able to separate talent, environment, and computing power well enough.

I do not think the answer is to search for the «pure» candidate, the one who has never used a machine. That would be artificial and probably counterproductive. Choosing a problem, guiding an automated exploration, understanding that a path is wrong, recognizing a good idea among a hundred mediocre proposals are all mathematical skills. What we would like to know is whether the candidate genuinely possesses them and can transfer them to a new situation.

Formalization helps certify a proof, not the talent of the person whose name is on it. It can tell us that an argument follows from the formalized assumptions; it does not tell us who had the idea, who understood where to look, or who will know how to change problems the following year. Even very transparent documentation of the tools used would not turn intellectual contribution into a measurable quantity. There is no subtraction of «value of the paper minus value of the tokens» that yields the value of the mathematician.

At this point chess almost entirely ceases to help us. We can organize a game or an examination without devices; we cannot compress three years of research into a supervised room. And I am not sure we should want to. Speed in solving an exercise before a committee is a real ability, but it is not the same as the persistence, intuition, culture, and judgment needed to keep a difficult problem alive for months.

We will probably have to return to looking more closely at people: discussing their work, asking unprepared questions, seeing how they react to a variation, working together when possible. Encounters both with and without tools could also be informative, provided they are not turned into a new evaluative fetish. No procedure will eliminate error: someone who explains a proof well is not necessarily the person who conceived it, and someone who hesitates when speaking may be an excellent researcher.

The 2022 Agreement on Reforming Research Assessment already moves in the direction of a more qualitative evaluation, entrusted to peer judgment and less dependent on automatic indicators [20]. This seems sensible to me, but also much more costly in terms of time and responsibility. Moreover, qualitative judgment can favor personal networks and prestigious institutions. If we really want to use it, we will have to create serious opportunities also for people coming from less visible environments.

Selection, in short, will also have a cost. Reading a small number of papers carefully, talking mathematics with candidates, and following them over time requires competent and available people. If the quantity of results produced increases while the time available to evaluate them decreases, the temptation will be exactly the opposite: to use automatic metrics and the reputation of the group. It would be ironic to economize on human judgment precisely when it becomes more valuable.

I do not have a recipe. Transparency, more equal access to tools, and less mechanical evaluations can help, but they do not remove uncertainty. There is also an almost circular problem: if we give more computation only to young people who already have the best results, we may end up rewarding those who already had access to the best computation. Some degree of equality at the starting point should, at least in part, precede selection.

This, however, is a problem of a living discipline, not a dead one. If we truly thought that mathematicians would no longer be needed in a few years, there would be little point in worrying so much about how to recognize the good ones. I do not know what the solution will be. I do know that stopping their training today because we fear they may be useless tomorrow would seem to me the riskiest choice of all.

11. Conclusions: the future of the profession

I do not know whether greed will really save mathematics, or mathematicians. The question mark in the title is there deliberately. The idea I have tried to discuss is much less solemn: if the most powerful computation remains expensive and those who control it want to profit

from it, total automation may turn out to be economically less convenient than its technical possibility suggests. It could be so. It could also not be so.

Much will depend on relative costs, which today we are unable to predict. For some tasks a machine will almost certainly be the best choice; for others a person may be; for very many, I suspect, it will be more efficient to have them work together. A good human idea can save millions of attempts, but a machine can find the idea that did not occur to us. I see no reason to turn this possible complementarity into a metaphysical war between human and artificial intelligence.

What convinces me least in the reactions of these days is the speed with which an extraordinary result is turned into a universal prophecy. It is a recurring temptation: as soon as a technology crosses a threshold, we immediately imagine the limiting case and treat it as the inevitable future. Perhaps this time the pessimists will be right; it would be arrogant to rule that out. But today it seems to me that we know too little about costs, access, incentives, and institutions to regard the obituary of mathematics as a scientific conclusion rather than an emotional reaction.

Nor does government intervention seem to me, in itself, a bleak prospect. It could produce control, bureaucracy, and new inequalities; it could also make available to many tools that no department could afford on its own. The difference will be made by the rules of access and by the choice to fund infrastructure and people together. I would not want the defense of mathematicians to be reduced to the hope that machines remain expensive. That would be a rather sad defense.

The possibility I find most interesting is almost the opposite of catastrophe: a mathematics in which we finally manage to tackle problems we do not touch today because they require too much time, too many calculations, or too much verification. A large automated campaign could open a door; simpler proofs, generalizations, counterexamples, and new theories could then follow. The history of a theorem does not end when the first proof appears. Often, for us mathematicians, that is precisely where it begins.

I am not even convinced that more results would mean less interesting work. It could mean the opposite. Every new class of examples, every new technique, every solved conjecture shifts the boundary and brings other questions into view. Certainly, some activities will become useless and some skills will lose value. This has already happened many times with symbolic computation, numerical computation, databases, and computers. This time the scale may be different, perhaps enormously different; but the conclusion «then there will be nothing left to understand» still seems to me much stronger than the premises.

If greed really does play a saving role, it will do so in a rather ironic sense: it will remind us that artificial intelligence is not magic and is not free. But I prefer an optimism that depends less on price lists. More powerful tools could make more productive the people who know how to ask good questions, distinguish an idea from a mirage, and build a theory around a result. If that is so, investing in mathematicians will not be an act of nostalgia; it will be a reasonable way of using machines better.

In the meantime, we have a piece of good fortune that we may look back on with surprise a few years from now: for a relatively modest personal cost, we can experiment with very powerful systems. We do not know how long this phase will last. Prices may fall and access

widen; the opposite may happen. In either case, it seems sensible to make use of the present. Not to produce more text or more tokens, but to learn, verify, try things we would not previously have tried, and leave behind something that does not expire with the subscription.

In the end, my most serious concern remains recruitment. If results become easier to produce and harder to attribute, how will we recognize a genuinely strong young mathematician? How will we distinguish talent from access to better tools, without penalizing precisely those who know how to use them well? I see no simple solution. And perhaps this is the least reassuring part of the whole discussion.

But here too, with some hesitation, I choose an optimistic reading. If the decisive problem is recognizing and training the best young people, that will mean that we still need them. Perhaps there will be fewer of them, perhaps they will do a different job, perhaps they will work every day with tools that seem exceptional to us today. I do not know. What seems premature to me is to confuse an enormous transformation with an ending. After Navier-Stokes we have very good reasons to be impressed, cautious, and even uneasy. I do not yet see good reasons to be in mourning.

23 September 2026

Carlo Mantegazza

A handwritten signature in dark ink, appearing to read 'Carlo Mantegazza', written in a cursive style.

Bibliographic references

Online sources consulted on September 13, 2026.

- [1] OpenAI. *On the Navier–Stokes Millennium Prize Problem*. September 8, 2026; update of September 10, 2026.
- [2] Clay Mathematics Institute. *Navier–Stokes Announcement*. September 11, 2026.
- [3] Fefferman, C. L. *Existence and Smoothness of the Navier–Stokes Equation*. Official description of the Millennium Problem, Clay Mathematics Institute.
- [4] Acemoglu, D., Restrepo, P. *Automation and New Tasks: How Technology Displaces and Reinstates Labor*. Journal of Economic Perspectives 33 (2), 2019, pp. 3-30. DOI: 10.1257/jep.33.2.3.
- [5] OpenAI. *Pricing*. OpenAI API documentation, pricing consulted on September 13, 2026.
- [6] Arrow, K. J. *Economic Welfare and the Allocation of Resources for Invention*. In *The Rate and Direction of Inventive Activity: Economic and Social Factors*, Princeton University Press, 1962, pp. 609-626.
- [7] Maslej, N. *AI Index 2025: State of AI in 10 Charts*. Stanford Institute for Human-Centered Artificial Intelligence, April 7, 2025.
- [8] Anthropic. *How we built our multi-agent research system*. June 13, 2025.
- [9] Anthropic. *Pricing*. Claude Platform Docs, pricing consulted on September 13, 2026.
- [10] International Energy Agency. *Energy and AI*. IEA, 2025, executive summary.
- [11] CERN. *CERN releases report on the feasibility of a possible Future Circular Collider*. March 31, 2025.
- [12] IBM. *Deep Blue*. IBM History.
- [13] Hassabis, D., Silver, D. *AlphaGo's next move*. Google DeepMind, May 27, 2017.
- [14] Silver, D., Hubert, T., Schrittwieser, J., Hassabis, D. *AlphaZero: Shedding new light on chess, shogi, and Go*. Google DeepMind, December 6, 2018.
- [15] FIDE. *Google announced as Title Sponsor of World Chess Championship 2024*. September 12, 2024.
- [16] FIDE. *FIDE Laws of Chess taking effect from 1 January 2023*. Arts. 11.2, 11.3, and 12.
- [17] Stockfish. *Stockfish: Strong open-source chess engine*. Official project website.
- [18] European Commission. *AI Continent Action Plan*. 2025, institutional factsheet updated May 7, 2025.
- [19] U.S. Department of Energy. *Energy Department Launches ‘Genesis Mission’ to Transform American Science and Innovation Through the AI Computing Revolution*. November 24, 2025.
- [20] *Agreement on Reforming Research Assessment*. July 20, 2022. Coalition for Advancing Research Assessment (CoARA).
- [21] OpenAI. *About ChatGPT Pro tiers*. OpenAI Help Center, individual plans and usage limits.
- [22] Anthropic. *What is the Max plan?* Claude Help Center, August 7, 2026.
- [23] Anthropic. *Model deprecations*. Claude Platform Docs, model lifecycle and retirement.