

Who is afraid of supermathematicians?

Mathematics, artificial intelligence, and the mathematician's profession

«My salary at the end of the month»

A colleague's answer, at the end of a seminar, when asked about the applications of his results.

It is not at all clear where the developments of these months will lead us. It does seem useful to me, however, to separate two questions that in recent weeks have continually been conflated: what will happen to mathematics, and what will happen to mathematicians. They are not at all the same question. Indeed, I have the impression that part of the confusion, and also some of the very strong emotional reactions, arises precisely from failing to distinguish them.

A preliminary point, so as not to mix different issues. I will not discuss here the recent accusations of plagiarism or appropriation made in connection with some results obtained with AI, particularly in the case of Navier-Stokes. For the argument that follows, I will assume that the results were obtained, verified, and attributed correctly. If that were not the case, the problem would be different from the one I want to discuss here.

The impression, at least for those who in recent months have seriously tried to work in mathematics with the most powerful models from OpenAI and Anthropic, is that two new players have entered the community. Or, to use a slightly provocative expression, two new mathematicians. Perhaps even two supermathematicians. Setting aside for a moment whether they are human or not, what matters is the kind of work they are able to do: solve problems, suggest ideas, combine distant tools, find errors, produce proofs, persist for hours along different paths. On a growing range of problems their abilities are already impressive and, in some cases, seem to compete with those of the best human mathematicians, even when the latter work in groups. I do not think it makes much sense, in general, to establish who is «better» between a great human mathematician and one of these machines. There are problems on which a machine is formidable, others on which it is still fragile, and still others in which an experienced mathematician recognizes within a few minutes that the machine is going in circles. For me, the point is different. The threshold has been crossed far enough to make plausible a scenario that until very recently seemed like science fiction: having on one's desk a mathematical collaborator who, on certain tasks, is better than we are and, above all, continues to improve.

In itself, this does not seem at all negative for mathematics. Supermathematicians existed before. Terence Tao is enormously better than I am and than many other mathematicians, yet the presence of Terence Tao has not brought mathematics to an end for everyone else. Mathematics has always been an ecosystem with enormous differences in talent. We work alongside people much stronger than ourselves, learn from them, pose different problems, find niches, collaborate, and teach. I do not see why the simple appearance of a more capable agent should, by itself, imply the end of everyone else's activity.

Tao's case is, among other things, particularly curious. Among leading mathematicians he was one of the early enthusiasts for AI, not in the sense of uncritical acceptance, but certainly one of the first to take it seriously as a research tool, to use it, and to argue that mathematical practice would have to adapt. For this reason, I am struck by the much more uneasy tone of his recent interventions. He has certainly not become anti-AI and continues to use it; but in September 2026 he was among the first signatories of the statement «A Severe Misalignment of AI in Mathematics», which is highly critical of the way companies are using the solution of mathematical problems as a measure of progress in their systems. Reading his statements over the past few weeks, one sometimes senses deep unease, not so much about the machines' mathematical ability as about the speed with which that ability may alter the values, practices, and institutions of the community. It is an interesting change of tone precisely because it comes from someone who embraced these technologies very early and very favorably.

Alongside Tao, Perelman comes to mind for almost opposite reasons. If one is looking for a symbol, Perelman is perhaps the last famous great mathematician «without AI»: not literally the last author of a major result obtained without these tools, of course, but one of the last iconic figures of a world in which an achievement such as the solution of the Poincaré conjecture could still be presented, at least in the collective imagination, as the «heroic» work of a single mathematician, although it built on Hamilton's earlier and fundamental results. It is also significant that Perelman was extremely harsh about the social dynamics of the mathematical community. He declined the Fields Medal and later the Clay prize; he explained that he felt alienated from the organized mathematical community and did not share its decisions and values. In his rare public statements he criticized conformity and tolerance of behavior he regarded as improper. Those criticisms concerned the system of prestige, recognition, and, ultimately, power within the community. Even before the current technological revolution, one of the greatest mathematicians of our time had already sharply separated mathematics from the «institutions that govern it». The difference, of course, is that Tao and Perelman are human exceptions, whereas an artificial system can be replicated, updated, and set to work in parallel. It is this scale, in addition to capability, that makes the transition entirely new and unpredictable. If tomorrow even stronger systems appeared, from different companies or developed in academic or government settings, we could have a plurality of artificial supermathematicians with different styles and strengths, and every human mathematician could use them as collaborators. The fact that someone, human or artificial, is better than I am does not automatically make what I do useless; the fact that this someone can be multiplied a thousandfold changes the picture much more seriously.

There is another point that may become important. Every mathematician is, in different proportions, a problem solver and a builder of structures, concepts, and theories. The proportions vary enormously from person to person. Riemann is perhaps the most immediate example of the second extreme: he introduced new objects and points of view that revolutionized entire parts of mathematics and created new ones. Erdős, simplifying just as brutally, better represents the other pole: a life made of problems, conjectures, lemmas, solutions, and collaborations, with an almost inexhaustible talent for seeing what was worth attacking and how to attack it. Naturally, neither can really be reduced to a single category: every mathematician carries a different mixture of the two components. For now, artificial intelligence seems to me to lie almost entirely on the problem-solving side. One asks it a question, or gives it a sufficiently well-defined problem, and it tries paths, combines techniques, checks cases, and searches for a proof. In this sense it can already be extraordinarily strong. It seems much less clear to me how well it can perform the other task: deciding that the existing language is the wrong one, inventing a new structure, building a theory that does not merely close a problem but changes the way we formulate many others. Today I would say, somewhat schematically, that AI is almost one hundred percent a problem solver. But I see no reason to turn this snapshot of the present into a prediction about the future. If this second ability were also to emerge, the change would be even deeper than the one we are observing now.

I will make an almost banal observation: if mathematics advances, in whatever form, it is difficult to regard that as negative for mathematics itself. If new theorems are proved, open problems solved, new structures built, and connections found that we had not seen before, mathematical knowledge grows. Whether the agent producing the decisive step has a biological brain or a neural network does not change the content of the result. A true theorem does not become less true because it was found by a machine. For centuries, in very different forms, sovereigns, states, universities, and societies have considered it useful to support and fund people who produce mathematics. Euler was called to Berlin by Frederick II of Prussia in 1741 as a «court mathematician»¹. This was done because the usefulness of mathematics to all the other sciences and to the development of new technology was understood, even though we almost never know in advance which mathematics will turn out to be useful. There is an amusing near-exception to this picture that I have always found entertaining, although it appears to be mainly cinematic. At the end of *Z* by Costa-Gavras, filmed while the Colonels' dictatorship was in power in Greece, the narrator lists the things banned by the regime, and «modern mathematics» appears among them. It would be a perfect counterexample to governments'

¹ What exactly did the «court mathematician» do? In Euler's case, the answer is: a little of everything. He did research, directed the mathematical section of the Academy, and was also consulted on scientific, technical, and practical questions.

confidence in mathematics, if it were true². Euler is also associated with a very famous, probably false anecdote that comes naturally to mind here, given the subject. During Denis Diderot's stay at the court of Catherine II, the philosopher is said to have spread his materialist and irreligious ideas with some enthusiasm. According to the story made famous in the nineteenth century by Augustus De Morgan, he was told that Euler possessed a mathematical proof of the existence of God. Before the court, Euler supposedly approached Diderot and, with complete seriousness, said: «Monsieur, $(a + b^n)/n = x$, donc Dieu existe; repondez!». Diderot, intimidated by the formula, is said to have been left speechless amid the laughter of those present. As a story, it is almost certainly too good to be true: Diderot knew mathematics and had written articles on it as well. The point of the scene is the authority of a formula that the interlocutor is unable, or does not feel able, to judge. It is not hard to see why this comes to mind today. If a machine much stronger than we are gives us an answer, perhaps accompanied by a formal certificate thousands of pages long or by a derivation that no human being will read in full, the problem will not only be knowing whether the answer is correct. It will be understanding what relationship we want to have with a mathematical truth that we can certify without truly possessing it in the human sense of the term.

This brings us to a question that interests me a great deal: what is a proof for? To a mathematician the question sounds almost absurd. Mathematics is not a collection of yes-or-no answers; it consists of proofs, structures, explanations, connections, ways of thinking. But do the applied sciences care about all of this in the same way? A physicist, an engineer, a chemist, or a biologist wants to know whether a conclusion follows from the model, under what assumptions, with what error, and within what limits. The proof can be extremely useful, but it is not necessarily the primary objective. Let us then imagine an extreme scenario, which is not the current situation: a machine reliable enough to be treated, for many purposes, as a black box. We give it a model and ask: is this conclusion true? Is this regime stable? Does this functional have a minimum? Does this procedure converge? Is there an estimate with these dependencies? The machine answers, perhaps also producing a formally verifiable certificate. For many applications that could be enough. A proof readable by a human being would become a second level: we would ask for it when we want to understand why the answer is what it is, not merely to know whether it is correct. Naturally, the word «infallible» would be out of place today. Models make mistakes, and in mathematics in particular they can make very convincing ones. But the conceptual point remains. If certification of correctness were separated from human understanding, one historical function of proof could be delegated. Not all of it, however: a proof does at least three different things. It certifies that a statement is true, explains why it is true, and, in the best cases, generates further mathematics. It introduces a method, shows which assumption really matters, suggests generalizations, and reveals a structure. It is likely that the first function will be fully automated much earlier than the other two.

In that case we might come to have true theorems whose complete proofs are too long, too technical, or simply too alien to be interesting to read. The mathematician might then ask the machine not «prove it to me» but «help me understand why it is true». These are two different requests, which until now we have almost always treated as one. And perhaps part of the mathematician's profession will shift precisely there: no longer only producing the proof, but extracting the explanation from it, recognizing what is new, turning a solution into theory. This leads to another possibility: part of our work could become mediation between the applied sciences and the machines that solve mathematical problems. The work would be to understand what a scientist is really asking, turn it into a sensible mathematical question, choose the right level of abstraction, check the assumptions, query the AI, interpret the answer, and explain which parts are robust and which depend on the model. In other words, a mathematician could become less the final solver and more the conceptual mediator. That does not seem to me a minor role at all. In many cases it is already the hardest part of our work: formulating the right problem is often more important than solving the wrong problem impeccably. Still, I would not place too much hope in this either. Nothing guarantees that an AI capable of doing mathematics will not also learn to talk perfectly with a biologist or an engineer, reconstruct

² As a historical fact, it does not hold up: studies of Greek schooling in those years show, on the contrary, that it was precisely under the junta that the reform of so-called «modern mathematics» was extended to secondary education. The image remains effective for the film, but this exception therefore appears to be false.

the implicit model, choose the assumptions, and explain the result better than we can. The role of interpreter may be a phase of the transition, not an eternal promise of professional survival.

It is precisely here, in my view, that we need to return to the initial distinction. For mathematics, the result matters. For mathematicians, who obtained it also matters. This is inevitable: human communities have to decide whom to hire, whom to promote, whom to listen to, and to whom to entrust resources. In academia, scientific recognition also turns into institutional power. Reputation matters in hiring competitions, invitations, committees, journals, and the distribution of funding. There is nothing particularly strange about this: every community has to organize itself. But the arrival of AI touches exactly the mechanism by which we have assigned that recognition. For a long time, with all its defects, the system has operated on the basis of a fairly strong correlation between individual scientific production and personal mathematical ability. A person writes papers; the papers are evaluated; some journals acquire prestige; that prestige reflects back on the authors; the authors obtain positions, funding, prizes, invitations, influence. No serious mathematician truly confuses the name of the journal with the quality of a theorem, but in practice the journal's name, the quantity and quality of the results, letters, and the reputation of the group have been used as signals of the quality of the person.

Artificial intelligence can break precisely this correlation and first of all raises a problem of assigning credit. If a substantial part of a result was obtained with the help of a machine, how much of the result do we attribute to the mathematician who signs the paper? If in a few hours I obtain a proof that on my own I would not have found in six months, what does my paper really say about my mathematical ability? How much of the result should be attributed to me and how much to the machine? Does skill in using the model count, the choice of problem, the ability to formulate the right question, to recognize the correct answer among many responses, to check and understand the proof? But how much does each of these things count? I see no simple answer, and above all I do not think it is enough to add a standard statement of the form «an AI system was used to...». That would be a useful form of transparency, but it would not turn an individual's intellectual contribution into a measurable quantity. And perhaps this is where one of the most traditional mechanisms of the mathematical community begins to crack: using the results produced as a more or less reliable indication of the ability of the person who produced them.

From this comes what seems to me the most concrete and urgent problem: recruitment. If two candidates present work of comparable quality and a substantial part of the intellectual work was done with artificial tools, how do we compare them? Prizes might also change meaning, or even be abolished altogether, and the journal system might enter a crisis at least in the function it currently serves as an indirect device for evaluating people. If producing a mathematical result becomes much less costly, the simple count of results, even refined by the reputation of the venue, says less and less about the person presenting them. Perhaps we will have to return to forms of evaluation that resemble a mathematical conversation more than bibliometrics: understand what a person can see, what questions they can formulate, whether they can distinguish a deep idea from a technically correct but sterile solution, whether they can explain a result, modify it, carry it elsewhere. This too, however, is easier said than done. An interview measures some abilities and misses others; someone who speaks well is not necessarily someone who thinks better; someone who hesitates before a committee may be an excellent researcher. There was no perfect system before, and no perfect one will suddenly appear now.

There is also the problem of access to the tools, which directly affects recruitment. It is not enough to say that two candidates have both «used AI»: what obviously matters is also which system they were able to use, with what limits, and with how much freedom to try. Today very powerful tools are accessible even to an individual researcher at relatively modest cost. If that remains true, AI could even reduce some inequalities. If instead truly important results require computing power that an individual or an ordinary department cannot afford, new inequalities will arise. In that case the scientific result will increasingly reflect the power of the infrastructure as well: what will count is not only the researchers' ideas and ability, but also the economic and computational resources available to them. I do not know which of the two trends will prevail, but for evaluating people the difference is enormous. A young person who can make a hundred times more attempts than another is not simply doing the same research with a slightly better tool. 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. For this reason too, the final product alone risks telling us less and less about the person presenting it.

This is also why some current reactions seem understandable to me, even if they do not convince me. Calls for «AI-free mathematics», bans, and attempts to construct artificially protected zones have a logic: they seek to preserve a ground on which individual performance once again becomes legible. But research is not a controlled experiment that can be isolated from the rest of the world. A paper takes shape over months or years, among readings, seminars, emails, calculations, other people's ideas, and failed attempts. I do not see how one could seriously certify that no step ever received artificial assistance. Frankly, I am not even sure that doing so would be desirable.

The problem becomes more delicate when we speak of so-called «useless» mathematics. For a long time we accepted a kind of black box called the mathematical community. Society funded research without demanding that every individual project have a visible application. The idea was: let mathematicians work; they will produce tools, theories, and structures, and some of all this, often in unpredictable ways, will become useful elsewhere. This trust allowed enormous areas of mathematics apparently remote from applications to survive and grow. Hardy is an almost too perfect example. He proudly defended «real» mathematics precisely as substantially useless and regarded number theory as one of the purest examples of that uselessness. Subsequent history, from cryptography onward, has shown how wrong he was and how dangerous it is to decide in advance what will never be useful. This does not prove that every abstract theory will find an application. It proves something more modest and perhaps more important: we are very bad at predicting in advance which parts of mathematics will become useful.

And this is where the remark by the colleague I quoted at the beginning, whose name I will not give, returns. At the end of a very abstract seminar, someone asked him what the applications of his work were. He calmly replied: «My salary at the end of the month.» The joke works because it is brazenly honest and because it touches on a real point. Not all mathematics has been produced with an immediate external justification, and that has not made it devoid of value. Some of it existed because someone found it interesting and a community decided that it was worth giving that person time to think about it.

What happens if, in the future, the applied sciences can query extremely powerful mathematical machines directly? Pressure may increase toward mathematics demanded by problems external to the discipline. In such a scenario, if a new structure is needed, the machine builds it; if a theorem is needed, it searches for or invents it. At that point it might become politically more difficult to justify entire human communities developing theories whose immediate use is visible to no one outside those communities. But this is not the only possible scenario. The exact opposite could happen: if producing and exploring mathematics becomes cheaper, we might be able to afford much more apparently useless mathematics, not much less. And there is an even simpler objection to a system that funded only the applied demand of the present: the machines themselves work within an enormous conceptual inheritance, built in large part without knowing what it would eventually be used for. A machine that is exceptional at solving formulated problems might need precisely that wealth of concepts, examples, theories, and languages born from curiosity. Ceasing to cultivate it could also impoverish the applied mathematics of the future.

Ultimately, the question is cultural and political before it is technical. For centuries the implicit pact has been roughly this: we give a certain number of people time to think, without continually demanding an immediate application, because overall we believe something useful, or at least something valuable, will come from it. The figure of the court mathematician is an old version of the same pact. Every now and then the sovereign asks for a fountain, a canal, or an artillery trajectory; otherwise he leaves the mathematician to do mathematics. The question is whether that pact will survive when the court mathematician can be a machine available twenty-four hours a day, one that does not ask for a chair and does not want a prize.

As for teaching, at least for now, the impact of AI seems less problematic to me. Teaching does not mean only transferring information. There is a component of presence, relationship, empathy, and intuition about the student in front of you. It also means understanding when someone has not understood but does not

want to say so, when an example needs to be changed, when to insist, and when to stop. It is possible that this part too will progressively be imitated by machines; I see no reason to declare it inaccessible in principle, but today the human element still has real weight. There is also a very simple point that we risk forgetting when we identify the value of a mathematician too quickly with the strength of their research: one can be a modest researcher and an excellent teacher. A mathematician who does not produce exceptional results can have enormous value in training students, transmitting ideas, and making mathematics understandable. This too is part of the profession and is by no means a secondary function. What we teach and how we teach it will probably change, however. It would make no sense to train students while pretending that these new tools do not exist. We will have to teach mathematics and, at the same time, teach how to use AI well, just as we have always taught students to write in LaTeX and to use MathSciNet, for example - things that perhaps, at least in part, will no longer be needed. We will have to teach how to ask a question, how to distrust an answer that is too smooth, how to ask for counterexamples, how to verify an argument, how to distinguish an idea from a sequence of formally plausible steps. Perhaps part of teaching will become much more oral. If a written exercise can be solved instantly by a machine, it becomes more interesting to ask the student to explain what is happening, react to a variation, correct an error introduced deliberately, or say why a certain assumption is needed. That would not necessarily be a bad thing.

There is also a value in mathematical education that survives even if a growing part of the production of results is delegated to machines. Studying mathematics trains rational and abstract reasoning, the ability to distinguish assumptions from conclusions, to build models, to recognize structures, and not to settle for a plausible answer. This is also why many mathematicians find work far outside mathematical research, universities, or research institutes: in scientific, technological, financial, industrial fields, or more generally in activities in which modeling ability and abstract thought matter. If the profession of the mathematical researcher changes, it by no means follows that this educational function of mathematics will lose its value.

All this also helps explain why the debate is so heated. Mathematics likes to present itself as the most impersonal of the sciences. Theorems do not depend on who proves them: they are either true or false. But mathematicians live within institutions, careers, friendships, rivalries, schools, prizes, journals, and seminars. We build part of our identity around the ability to understand and produce mathematics. Seeing a machine appear that can do in a few minutes something on which a person has built twenty years of expertise is not psychologically neutral: it can be emotionally dramatic. The most abstract of sciences reveals itself, without disguise, in its deeply human and social reality. For this reason I would not dismiss all fears as hysteria. Some are rational and concern real problems: jobs, access to tools, recruitment, power, funding for pure research. But I would not turn these fears into a requiem for mathematics either. Mathematics could be doing very well, perhaps better than ever, while mathematicians go through a serious crisis of identity and institutions. The two things can happen at the same time.

Finally, there is one variable that could determine much more than we imagine today: how much it will really cost to use the most powerful systems. Here I would distinguish at least three levels. There is the personal cost: if a mathematician can have, with an ordinary subscription, an extraordinary artificial collaborator, the transformation will be rapid and widespread. There is the departmental cost: if seriously attacking a problem requires thousands of agents, days of computation, or special tools, then AI resembles infrastructure more than a subscription, and not every group will be able to use it in the same way. And finally there is the social cost: energy, hardware, public or private investment, environmental impact, and, inevitably, resources that are not spent elsewhere. A technology may be capable of doing something and yet it may not make economic sense to use it for every problem. These three scenarios lead to very different futures. If frontier computing power becomes cheap, every mathematician could work daily with one or more artificial super-collaborators, and the transformation would rapidly affect the entire discipline. If it remains expensive, mathematics could instead depend increasingly on selective access, priorities, consortia, and differences among departments or countries. If the social cost were very high, a collective choice would inevitably enter the picture: not only «can we do it?» but «on which problems is it worth spending these resources?» For this reason, besides asking how intelligent AI systems will be in five years, perhaps we should also ask how much it will cost to put them genuinely to work.

I do not know how the profession, recruitment, journals, prizes, authorship, or teaching will evolve. I do not even know how long it will remain meaningful to distinguish between «mathematics done by humans» and «mathematics done by machines», because it is likely that a growing part of real mathematics will become a hybrid production. The interesting problem will then be to understand what, within that hybrid, we want to continue calling human mathematical activity, and why. Perhaps we will continue proving theorems simply because we enjoy doing so, because understanding an argument, finding an elegant proof, or managing to explain it to someone else are activities that have value even when the answer is already available. Perhaps the aesthetic, cultural, and even playful dimension of mathematics will become more explicit precisely if some of the professional reasons for doing it disappear. Perhaps we will discover that understanding a proof is a different activity from possessing the answer. Or the change may be more radical than we can now imagine. Chess and Go provide, at most, a very partial precedent: machines have long surpassed the best human beings without making human play meaningless. The analogy with mathematics has obvious limits and I would not want to push it further, but it at least reminds us that the technical superiority of a machine does not by itself determine the value of a human activity.

For now it seems premature to me to defend ourselves by banning the tools or pretending they do not exist. It also seems premature to celebrate the end of a profession when we do not yet know how it will change. The most interesting question is not how to preserve intact the system we had, nor how to beat the machines at their own game. It is to understand what mathematics, what mathematicians, and what mathematical community we want to build now that, to use the words from the presentation of Nello Cristianini's essay *Machina Sapiens*, «perhaps we are no longer alone». For the moment, perhaps it is best to settle for Turing's caution: in 1950 he ended his article on thinking machines by writing, «We can only see a short distance ahead, but we can see plenty there that needs to be done». For mathematicians too.

25 September 2026

Carlo Mantegazza

References and sources

Euler in Berlin: Leonhard Euler was called to Berlin by Frederick II in 1741 and worked at the Academy of Sciences, dealing with practical and institutional questions as well as research.

Euler and Diderot: the best-known version of the anecdote about the «algebraic proof» of the existence of God appears in Augustus De Morgan, *A Budget of Paradoxes*, vol. II, drawing on an account by Dieudonne Thiebault. Dirk J. Struik, *A Concise History of Mathematics*, 3rd rev. ed., Dover, 1967, p. 129, regards it as an example of bad historical anecdote and notes that Diderot knew mathematics.

Terence Tao and AI in mathematics: T. Tao, *Mathematics in the age of AI*, public lecture at ICM 2026 and arXiv:2608.16753; see also *A Severe Misalignment of AI in Mathematics*, September 11, 2026, of which Tao was one of the 25 initial signatories. As early as 2023, Tao was publicly arguing that AI tools would profoundly transform mathematical practice.

Perelman and the mathematical community: the official account of ICM 2006 in the *Notices of the AMS* reports that Perelman felt isolated from the mathematical community and did not believe his values were represented there; S. Nasar and D. Gruber, *Manifold Destiny*, *The New Yorker*, August 28, 2006, report his criticisms of conformity and the community's ethical standards. In 2010, when declining the Clay prize, Perelman also indicated his disagreement with the organized mathematical community and its decisions.

Hardy: G. H. Hardy, *A Mathematician's Apology*, Cambridge University Press, 1940, for the defense of pure mathematics and the theme of its apparent uselessness.

«Modern mathematics» under the Greek Colonels: the phrase appears in the final enumeration in the film *Z* (Costa-Gavras, 1969). It does not, however, correspond to a general ban on «modern mathematics»: Dionysios Baltzis, in «Hands off Euclid!»: The Ambiguous Reception of the «Modern Mathematics» Reform in Greece (2023), pp. 547-552, shows that under the junta the reform was in fact extended to secondary education.

Cristianini: N. Cristianini, *Machina Sapiens. L'algoritmo che ci ha rubato il segreto della conoscenza*, il Mulino, 2024; for the phrase «Perhaps we are no longer alone», see the publisher's presentation of the volume on the il Mulino website.

Turing: A. M. Turing, *Computing Machinery and Intelligence*, *Mind*, 59 (1950), no. 236, pp. 433-460, in particular p. 460; DOI: 10.1093/mind/LIX.236.433.