1 Theorem, Truth, Agency, Measure

1

Suppose a theorem is proven today. The belief prior to the AI era would be that the action behind the proven theorem is that of a human. How this belief arise in the first place? Surely because, for millennia, the culture, the folklores, and the stories, had collectively carried this belief, ultimately granting it to the humans of the current era. For the sake of argument, let us keep the belief aside, and look at what has happened today, as when the theorem is proven. (We keep aside, for now, the question of where the truth came from, and refer the readers to the discussion by Roger Penrose in The *Emperor's New Mind* (Vol. 19, 1989) Oxford: Oxford University Press on Platonism versus Intuitionism, p. 112).

1.1

What's relevant to the forthcoming discussion is that the truth born today must have had an agency behind it. By the current state of affairs (post Navier-Stokes fiasco), the agency is exercised by either: a human; or a machine; or them as a team. What carries more importance: an agent; agency; or the theorem? It depends on what you mean by important, and what's important to one could be not so to the other. Regardless, the truth, hidden inside the statement of the theorem itself, born today can be safely said to have a primordial importance, for its immaterial nature that correlates, aligns, and materialises with the rest (agent, agency, theorem, and discourse around it). Truth is primordial.

1.2

The theorem stating the truth can be assigned a measure for us to start the discourse. The first measure is how non-trivial its proof is to our world: some theorems appear longer and winding, while others appear shorter and sleeker to us. The second is the extent to which the agent has delved deeply into themselves to excavate that "reproducible" [c.f., Venkatesh, A.] truth, scaffolded in the statement and shown to us (the others). There must be some process that the agent has gone through. What exactly is this process and to what extent the agent has carried themselves through it? Would we want to measure it or talk about the process, or leave it as is and just care about the theorem that has appeared in front of us?

1.3

Let us take the first measure. The non-triviality of a proof is communicable, very intrinsic and personal to the mathematical community; the longer one has spent thinking and calculating, the larger the number of pages that the proof occupies. One might ask, what exactly is a page, and is the size of the page a good metric for the measure we are interested in? Think about how an A4 page versus a microfiche page can be compared quantitatively through a ratio, but certainly two agents who wrote in A4 and microfiche respectively cannot be compared quantitatively through such a ratio. But, why not? Isn't it easy to say that agent A wrote a proof for statement X in 288 pages while agent B wrote a proof for statement X in one microfiche page (which is equal to 576 A4 pages), and thus agent B is 1.25 times better? Agent B could be better in efficiency, but how about quality? Quality is a soft metric, that we could be more concerned with, because it's much more intrinsic to the community. Pictorially and metaphorically, like how a hard shell of an egg hides the soft interior of its yolk, the hard metrics of an agent (which are easy to communicate and have discourse about) are often covered by soft metrics of an agent's quality. This metric is told repeatedly, through writings, folklores, stories, books, poems, and all sorts of social, human, essential things, over not just a couple of days, but months, years, decades, centuries, and even millennia. On the discussion of the preservation of mathematical knowledge, we refer the interested reader to *Gathering, Preserving, Transmitting: Mathematical Heritage through Encyclopedias*, 1700–2000 (2026) Birkhäuser/Springer Nature Switzerland.

2 Economy of Measure

2.1

The first measure, seemingly hard, and the second measure, seemingly soft, in totality has created an economy of mathematics: that of suppliers and that of demanders. The suppliers from the book (c.f., Erdos, Proofs from THE BOOK) have infinite problems to supply (pose) to us, and the demanders demand: (a) second measure – an infinitely rich, inner experience, and (b) first measure – a suitable number of refillable pens and paper, for a perfect match to happen. For centuries if not millennia, such a match has repeatedly happened. Except, in the cases of arbitrage, which we defer to Section 2.4.

2.2

The action of machine can either replace or supplement the action of human; and the way leading to the match. The first case of replacement is when the cost of machines producing theorems is lower than the infinite cost (pricelessness) of the book to exist and pose problems at the first place (met by cheap, and getting more cheaper, the machine tokens). The second case of supplementation is when the cost of machine-human producing theorems is close to or exactly matches the infinitary cost of the book. But why should the book always have an infinitary quality, and how could the cost of producing be infinite? There could very well be a portion or a version or an adaptation of the book, with a fixed, finite cost, demanding machine-humans to produce theorems at a cost (and effort) matching the book's expectation.

2.3

What is the cost of a machine-human collaboration? What is the cost of an institution keeping the book (or perhaps, its adaptation)? A complicated question like this is, in its essence, the question of where the supply costs and demand costs, structurally, fall. To whom does the machine-human collaboration belong to? Whom does the adaptation of the book (if not the book) belong to? Regardless of the institution (be it the Bay area or a university or an archivist's corner), one thing which is true is that the machine-human collaboration has a cost of the supply of its theorems that can embed nicely, and nearly or perfectly match, the cost structure of the demand of the institution. But what about the measure of the human-machine collaboration? If the institution holds a finite adaptation of the book, with its cost well bounded below its endowment, the institution cannot allow a potential infinitary, rich experience of the human-machine collaboration to be matched. It is certainly the case that a human-machine collaboration could be finitary and that is only the case when a match will happen.

2.4

Mathematics can have some portions which could be scaled up and industrialised (proposals like CERN), in the sense of invoking division of labour and building organisations and institutes achieving massive-scale, crowdsourced, citizen science style projects (c.f., Tao, T.). Because of finite cost and mismatch between expectations and demands, small amount of arbitrage are possible. However, mathematics as a science can be built as a well-coordinated discipline of the kind other fields already are, and thus, reduce these chances of arbitrage, and help in matching the suppliers with the demanders. Let us for now focus only on a singular human-machine, atomic case.

2.5

A human-machine experience, of an infinitary kind, can go to a metamathematical wilderness, to explore the mathematical questions of the book, and thus, aim to achieve an infinitary experience of a rich kind, that previously solely humans (or collaboration between humans, c.f. Polymath Project) could experience (and an ecosystem, like, inter-institutional collaboration, digital communities, inter-cultural collaborations would grant them). Let us start imagining a proto-ecosystem, for the sake of an argument – a human and a machine sitting in a Norwegian hut (c.f., Ludwig Wittgenstein), and exploring with all the conviction, judgement, passion, and love, the vast, infinite mathematical space of the book, and aiming to experience a collaboration that can match the extent of the book. For this arrangement to be effective economically, we want this merger to happen at a just cost: the infinite cost of the book's demand, matched by the infinite cost of the human-machine experience required to meet that demand.

3 A human-machine experience

3.1

In this metamathematical hut, the human-machine is located far and isolated with just enough resources to survive and work. Let us invoke Penrose's model of interconnected worlds (c.f., this blog), where the mathematical, the physical, and the subjective worlds communicate to each other through ways shown in the figure below (left). Placing the human-machine in Penrose's model means decomposing what's happening into the dynamic represented in each of these three worlds. The machine finds itself in the hut, and since it belongs to 'Nature', the machine can be said to exist in the physical world. This scenario is helpful to us to explore the question posed in 2.4. Clearly, the machine on its own would be incomplete, yet still capable of explaining some aspects of the reality – as other physical models are (capable yet not complete). The issue that sits at the core now is: when would the machine, while working with the human as an agentic team, take on the task of exploring the mathematical world and the infinity of mathematical richness, and resolve our question.

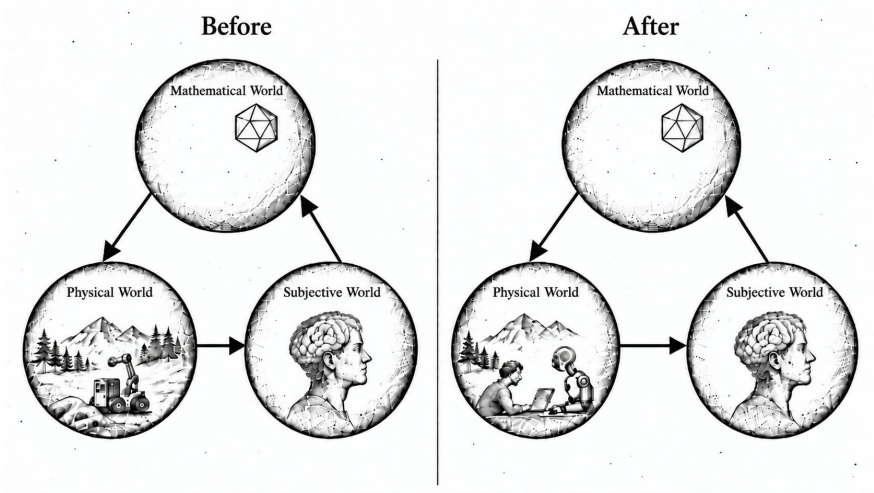

Figure 1: (left) Penrose's three interconnected worlds and (right) its modification [generated by GPT-5.6 Sol]

3.2

For centuries before the machines came to the fore, the human was investigating questions of a mathematical kind, to understand the physical world, and in the process experience a state of inner richness. A great portion of this process (ref. 1.2) is undergoing a transformation. The machine, previously mechanical, is now intelligent, and can identify truths, far more advanced than 1+1=2, from the mathematical world, and generate, liven them up as statements in front of us. But is this process of machine-generation complete and infinite in the way human is used to and would like it to have? It is certainly the case that human would like the machine to shine some understanding and judgement (the version of that for humans would be as, narrowed-down, pedantic, slow steps, that permit an infinitary richness of experience) and the chances of this are very low, because, the "non-responsiveness of the machine" (c.f., not embedded in an order of interpretation and meaning, Friedman 2026 p.10), gives rise to anxiety. If not, then machine, being part of the physical world, would prompt (in fact, provoke) humans to go and find it. A machine with its vast compute, memory, skills (and coarse-grained judgement and cleverness) of a finite kind can indeed out-compute, out-calculate, out-memorise humans; though, the humans, being stuck inside their metamathematical hut, and decoupled from the physical world that the machine inhabits (ref. figure above - right), and focused deeply on their own, rich, inner, subjective experience of the mathematical world it inhabits, would take on those machine-generated slop and attempt to build infinite, human experience of the kind necessary to supply the just explanation and thereby, match the cost that the book has.

4 Back out of the hut with Peace

Once these infinite costs are demanded, and no arbitrage exists, have we arrived at an agreement? What is next? Before we try to jump quickly ahead, let us ask, what was it like to be anxious about the non-responsive machine? The Other did not care to say anything beyond what its symbols referred to. The symbols – living in a formal, symbolic space – captured tiny fragments from many mathematical islands we are familiar with. If each of those islands represented a key emotion distributed around other emotions, it is natural that the machine-symbols, if nothing else, have a property of invoking and inciting many such emotions within us, each specific to that mathematical island. Isn't it the case that we compartmentalise our emotions to different areas of enquiry, and keep, for example, our beloved island separate from a neutral island? A commutative geometer could, for example, be fearful about the island of noncommutative geometry for its lack of exiting the space feature, which leads to a general lack of theorems. Or a category theoretician could be dissatisfied with set theory and its foundations. A machine, belonging to the physical world, and living in a formal, symbolic space, with its ingenious capability to incite and invoke a variety of human emotions inside a human living in the metamathematical hut, seeking an infinite, rich, inner-human experience, would seek to absorb and acknowledge the plethora of such emotions. An ideal scenario is a state of peace: a calm, satisfactory state of mind in which the distinction between the machine-physical and human-mathematical worlds is acknowledged (c.f., Penrose's three interconnected worlds), while the emotions that the symbolic space of machines tends to invoke are set aside. After this rigorous exercise, what is the human left with? A state, beyond reason, of love, joy, and an ultimate peace, which doesn't need to be constant and permanent. It could arrive on some special occasions, and remind the human that mathematics has had a role in helping the human – sense, absorb, reflect, and forget the issues the physical world tends to have. And instead, in an AI era, the human can focus on using machine-knowledge to do the same again: to acknowledge and reflect what's around, but leave that aside, to continue seeking – and occasionally - experiencing the peace again. Collectively, all of us doing that would help ourselves *make peace with mathematics*.

Acknowledgement

The author is hugely indebted to Connor E. Hayes, Faculty of Divinity, University of Cambridge and Dr Pablo H. Haerle, Department of Philosophy, NYU, for important comments and clarifications on the use of themes such as agency, intrinsicness, mathematics of science, and for grammatical and tonal edits.

AI disclaimer

The author wrote this essay independently, followed by several rounds of discussion and revisions with colleagues mentioned above on the key philosophical terms. The final draft was polished using GPT-5.6 Sol (medium) solely for grammatical corrections.