FROM "COSMIC ENERGY" TO INSUFFICIENT EVIDENCE: HOW I GRADUALLY LEARNED TO GIVE ABSTRACTION A TYPE SYSTEM 🧠📚

AN ESSAY ON WRITING, MATHEMATICS, SCIENCE, AND UPDATABILITY

DORIAN VERLAINE

Abstract. This essay traces how training in writing, mathematics, and physics gradually turned boundaryless abstraction into models that can be defined, expanded, inspected, and updated. Beginning with the small debugger "What does that mean, concretely?", it examines argument structure, premises and evidence, experience and authority, falsifiability, communication across different backgrounds, and when to stop a conversation with no convergence condition. Its central claim is not that abstraction or intuition should be rejected, but that abstraction should have a schema, intuition should carry a confidence level, and every model should retain an update path.

There is a certain kind of writing that creates a very peculiar feeling: it says a great deal, every sentence seems vaguely wise, and yet by the end it is difficult to answer the most basic question of all: **what did it actually say?** Phrases like "raise your energy," "return to your true self," "understand the frequency of the universe," "complete your life lessons," or "enter a higher dimension of awareness." They often sound smooth and polished, even carrying that texture of "this seems to touch something profound." But ask one simple question—"What does that mean, concretely?"—and the whole system may start dropping packets. 🤡

This conversation originally began with a self-help book. I was mostly reading it out of curiosity, just to see what kind of creature you get when you stuff "worldly wisdom, psychological growth, systems thinking, Eastern philosophy, and cosmic energy" into the same volume. But as the discussion went on, I gradually realized that the really interesting question was not the book itself. It was something broader: **why do people package wisdom into structures that become increasingly abstract, increasingly difficult to inspect, and sometimes so vague that even the speaker may not fully know what they are saying? And how did I gradually find my own way out of that mode of thinking?**

1. Abstraction itself is not the problem. Abstraction without boundaries is.

I used to slip very easily into a kind of "boundaryless abstraction." After seeing a few concrete things, my brain would immediately start compressing:

> Friendship → relationships → power → structure → systems → existence → the universe 🤡

That compression is extremely seductive, because it creates the feeling that "I am seeing a shared structure other people have missed." The problem appears when abstraction develops faster than the mechanisms used to verify it. A lot of ideas are left as nothing more than a vague sense of the whole, without sufficiently clear definitions, examples, counterexamples, or intermediate reasoning. You end up thinking, "I obviously came up with something huge, so why can't anyone understand

Date: 2026-09-11.
2020 *Mathematics Subject Classification.* Epistemology, Reasoning & Technical Metaphor (Essay).
Key words and phrases. epistemology, abstraction, type-systems, reasoning, communication.

me?" Yet when asked to identify exactly what the other person failed to understand, you cannot explain it either.

This is where a dangerous substitution can happen: treating "other people cannot reconstruct my idea" as "other people are not intelligent enough to understand it." Quite often, a more accurate description would be: **there is a highly compressed semantic blob in my head, but I have not yet factorized it into a path another person can actually follow.**

Good abstraction does not escape from the concrete. It imposes stricter constraints on it. Mathematics is abstract not because it is free to float wherever it likes, but because it has a type system. If you say "vector space," a set of axioms must hold. If you say "linear map," addition and scalar multiplication have to be preserved. If you say "limit," you cannot stay forever at the intuitive level of "getting closer and closer"; eventually you have to return to ε and δ. Here, abstraction is not fog. It is a compression format, and it can be decompressed.

The real problem with many "mystical" expressions is not that they are abstract, but that they are **abstract without constraints**. If someone says, "Your energy is wrong right now," and when asked what that means they answer, "It means your state is bad," then they have merely dressed an ordinary concept in a cosmic skin. If they go one step further and say, "People with the same frequency attract one another, so the universe sends people of the same frequency into your life," things become even more interesting. The "frequency" that began as a metaphor has quietly been upgraded into a mechanism. The metaphor has rooted the whole machine. 💀

So I now have a very simple debugger that I like enormously:

> **What does that mean, concretely?**

"Raise your energy."

Concretely?

"Maintain a regular sleep schedule, exercise, and reduce relationships that leave you chronically exhausted."

Good. Now the sentence has finally landed.

If the answer is still "bring your frequency back to a higher state," then ask again:

> **What does that mean, concretely?**

If ten recursive calls still fail to touch observable reality, then at minimum we can attach a label:

```
meaning: unresolved
confidence: low
```

There is another distinction that took me some time to notice: **compression and vagueness are not the same thing.** Good abstraction really can compress a huge amount of detail into a short expression, but that compression should be reversible. When necessary, I should be able to expand it again through definitions, reasoning, and examples. Mathematical notation can be extraordinarily compressed without being vague; a theorem may fit on one line, but every symbol behind it has a fixed interface. By contrast, terms such as "life energy" or "higher-dimensional awareness" can also look highly compressed, yet often have no stable decompression algorithm. They are not compression so much as lowering the resolution of an image and then claiming the picture has become deeper. 🤡

So when I encounter an abstraction that "sounds profound," I now ask another question: **can it be expanded?** If a concept can unfold from one sentence into definitions, conditions, examples, and counterexamples, and then be compressed back into the original sentence, it probably does carry information. If expanding it merely produces another cluster of equally vague terms, then what it carries may mostly be atmosphere.

2. IELTS writing unexpectedly rewired the way I think

What gradually made me aware of these problems was not mathematics alone. IELTS writing affected my entire style of expression far more than I expected. On the surface, it trains English. In practice, it also works like a rather brutal cognitive hygiene system: what exactly are you claiming? What is the central point of this paragraph? Is this sentence related to the previous one by causation, contrast, addition, or example? You wrote `therefore`—but does the previous statement actually imply the next one? Is your example supporting the argument, or merely performing next to it?

After enough practice, your brain grows a tiny reviewer:

> "This sentence is beautiful."
> "And?"
> "..."
> "Delete it." 🤡

IELTS also forced me to build a "reader model." My older style of expression often assumed something like this: these concepts are obviously connected in my head, so surely you can see the connection too. But the reader does not have SSH access to my brain. I have to send the intermediate packets. I have to make the implicit relationships explicit. Why does A lead to B? Under what conditions does B lead to C? Which step is correlation rather than causation? Which example is merely an illustration rather than evidence? None of that can be left for the reader to guess.

Curiously, this later began to echo what I was learning in mathematics. Mathematics says: "Do not skip steps." Writing says: "Do not make the reader guess what you skipped." One demands proof structure, the other argument structure, and both attack exactly the same sentence:

> "I feel there is some deep connection between them."

Teacher: What connection?
Me: It's just, you know, that kind of...
Teacher:

```
Band 6.0
```

🤡

Another major change was that I increasingly stopped believing that "big words" automatically equal depth. `profound`, `fundamental`, `essence`, `humanity`, `civilisation`—none of these words is forbidden. But the broader the scope of a sentence, the greater the burden of proof it has to carry. "Social media has destroyed authentic human connection" sounds dramatic, but "high-frequency use of short-form video may reduce the ability to sustain attention for long periods in some contexts" is much less explosive and much more discussable, testable, and falsifiable. **Being willing to shrink a sentence in exchange for verifiability is a very important kind of maturity.**

IELTS also left me with another durable distinction: "I have many ideas" is not the same thing as "I have an argument." The former only requires many nodes in your head. The latter requires edges, and those edges have to point in coherent directions. A lot of apparently information-dense writing is really just ten related words placed in the same paragraph. The reader senses that they belong to the same cloud, but cannot tell which edge is causal, which is evidential, and which is merely associative. Clear writing, by contrast, is more like making the graph explicit: this is the claim, this is the reason, here are the conditions, and this example only supports one specific part. From that point on, I became increasingly dissatisfied with "roughly speaking, that's what I mean."

3. LADR, Spivak, Stewart, Feynman: abstraction begins to acquire a schema

After working with LADR, Spivak, Stewart, and even *The Feynman Lectures on Physics*, this change became much more obvious. One of their most important effects was not merely that "I learned certain knowledge," but that they made me feel very strongly, for the first time, that serious abstraction has boundaries.

Spivak will not let me replace the rigorous definition of a limit forever with "getting closer and closer." LADR will not let me escape with "linearity means preserving some kind of structure"; it forces me to say which structure, which operations are preserved, and what exact properties must hold. Feynman's writing is highly intuitive, but he repeatedly pulls concepts back toward physical situations, dimensions, and observable consequences. All of them are doing the same thing: tying weights to abstraction so it cannot float off into the stratosphere.

I increasingly like to describe the transition this way:

```
Before:

idea: any

Later:

definition -> proposition -> argument/proof -> example -> counterexample
```

Abstraction finally has a type system. So now, when I see a sentence like "True wisdom is to understand the essence of life from a higher dimension," my brain is much less likely to produce 🤯 first. Instead, it produces:

```
define("wisdom")
define("dimension")
define("understand")
define("essence")
```

If all four are undefined, then at least for the moment it is not "deep." It simply has no schema.

The Feynman Lectures on Physics gave me another interesting reminder. Near the beginning of Volume I, Chapter 1, Feynman writes a very short sentence about science: **"the sole test of its validity is experiment."**[1] Of course, this cannot simply be applied to every philosophical, ethical, or value question. But for claims that say something about "how the world actually works," it is an excellent firewall. An idea can come from anywhere: intuition, metaphor, dreams, historical experience, even a beautifully constructed story. But once it begins to claim an observable causal relationship, it must be prepared to reconcile itself with the world.

In the same chapter, Feynman also presents a beautiful thought experiment: if all scientific knowledge were destroyed and only one sentence could be preserved, he would choose the atomic hypothesis—the idea that the world is made of moving, interacting atoms.[2] I like this passage not merely because it discusses atoms, but because it demonstrates a form of high-density expression that is almost the opposite of "mystical" language: one sentence can be extremely abstract and extremely compressed, yet expand downward into the properties of water, temperature, pressure, the structure of matter, and a vast number of observable phenomena. **A genuinely powerful abstraction does not make you stop asking questions. It gives you more questions that can be asked.**

---

[1] Richard P. Feynman, Robert B. Leighton, Matthew Sands, *The Feynman Lectures on Physics*, Vol. I, Ch. 1, "Atoms in Motion". Caltech online edition: The Feynman Lectures on Physics.

[2] Ibid. In §1–2, Feynman introduces the atomic hypothesis through the thought experiment of preserving only one sentence of scientific knowledge, and uses it to show how a highly compressed physical idea can unfold into a large number of concrete consequences.

That gradually gave me another criterion: if a "deep" statement becomes more concrete as you explain it, produces more predictions, and acquires sharper boundaries, then perhaps it really is compressing knowledge. If explaining it merely produces more "energy, frequency, awareness, universe, oneness," then semantic recursion may be calling itself. 💀

4. But then comes the problem: communicating with people who do not share the same training is exhausting

This change has side effects. Once I became used to distinguishing concepts more carefully, checking intermediate steps, caring about evidence, and noticing abstraction loss, conversations with people who have not gone through similar training sometimes became very expensive.

For example, I often use computing terminology as analogy. I might say, "This looks more like a cache miss than lower ability." To me, that one sentence already compresses a pile of prerequisites: what cache means, how hits differ from misses, why the first recomputation is expensive, why an existing model can dramatically reduce future processing cost. But if the other person lacks that background, then I have not really spoken one sentence. I have secretly depended on a 4 GB shared library. 🫠

Communication then becomes a real-time runtime-detection problem: does this person know the term? Do they realize I am using it analogically rather than literally? Is their nod genuine understanding, or merely a polite ACK? Are they rejecting my conclusion, or did they never receive the premise? If they do not share the same background, I have to change the encoding and invent the explanation again.

I increasingly think that a more efficient approach is not "explain everything in full every time," but to use a layered protocol. The same idea can have several versions:

```
level 0: Understanding something for the first time can simply be slower.
level 1: Many apparent differences come from different existing
         mental models, not necessarily different abilities.
level 2: I would compare it to a cache miss / cache hit.
level 3: Then we can discuss representation, reconstruction cost, and transfer.
```

Send the minimum version first, then decide whether to expand based on the response. Do not begin with `level 3 multicast` across the entire network. 🤣

More importantly, do not merely ask, "Do you understand?" because the answer is often just:

```
ACK
```

A better checksum is: "What did you understand me to mean?" or "What do you think the parts of that analogy correspond to?" This is not an exam. It is a corruption check on the payload.

5. The truly difficult problem is not different background knowledge, but different epistemic protocols

Differences in background knowledge are actually manageable. At worst, you perform a few extra rounds of translation. The more exhausting case is when the other person has deeply rooted beliefs and almost never questions them.

At that point, the difference is no longer merely "we know different things." It becomes "we disagree on what counts as a reason." I may ask: what is the evidence? Is there a counterexample? Under what assumptions does this conclusion hold? Has anyone tested it? What were the results? How much information was lost when moving from a concrete result to an abstract conclusion?

The other person may answer: Professor So-and-so says the same thing; everyone has always done it this way; I have lived this many years and this is my experience; you are still young and inexperienced, so your view is unreliable.

One distinction matters here. "A professor also says this" is not automatically invalid. Relevant expert judgment can absolutely count as evidence. The question is how much weight it deserves.

Is this person an expert in the relevant field? Is the view itself supported by evidence? Is there a stable expert consensus? The rule should not be:

```
Professor said it -> true
```

Similarly, "you lack experience" can at most rationally reduce the credibility of some judgments that depend heavily on tacit knowledge. It cannot directly imply "your argument is wrong." Someone who has never flown an aircraft is indeed unlikely to describe the handling characteristics of a 737 reliably. But if that person points out that your fuel calculation is missing a factor, replying "You're not even a captain" does not fix the arithmetic. 💀

Pride makes things even worse. Someone may already suspect that their view is partly wrong, or entirely wrong, and still refuse to update because admitting error would damage their identity or status. Reasoning then shifts from "which model is closer to reality?" to "how do I protect my existing conclusion?" People begin selectively collecting favorable material, exaggerating experiences that support them, downplaying counterexamples, using authority as a substitute for argument, or attacking the other person's credentials.

That is what makes motivated reasoning so exhausting: **reasoning stops being a tool for revising conclusions and becomes a tool for protecting them.**

"Experience" is especially easy to misuse here, because experience obviously has value. Some knowledge can only be acquired through long practice. But there is still a long road between "I have seen this many times" and "this is a universal rule." Experience can suffer from selection bias, may cover only a particular period, region, or social group, and can be made artificially neat by the way memory is reconstructed. The danger is not using experience. The danger is treating experience as an uninspectable root certificate: once someone says "I have lived for decades," all downstream inference is automatically trusted. Epistemically, that is a very expensive privilege. 🤡

So I now prefer to treat "experience" as a prior or as additional evidence rather than as the final judge. Experience can raise the initial plausibility of a judgment, but it can still be revised by better data, larger samples, more direct measurement, or a more complete mechanistic explanation. Conversely, pure theory should not automatically defeat practice merely because its formal structure is elegant. The useful question remains: **which side offers a model that explains the phenomenon better and survives counterexamples more successfully?**

6. Before the conclusion comes another layer: what premises have we accepted?

Eventually I realized there is another layer that is even easier to ignore than "is the evidence sufficient?" or "is the reasoning valid?": **premises.**

Many arguments look like disagreements over a conclusion, when in fact the two sides were never standing on the same floor in the first place. A's reasoning may be perfectly valid, and B's reasoning may also be perfectly valid, but if they accept different premises, they can reliably arrive at different answers. At that point, continuing to throw more examples at one another at the level of conclusions usually only expands the battlefield.

The simplest form looks something like this:

```
Premise 1
Premise 2
Premise 3
    ↓
Conclusion
```

If the conclusion is wrong, there are at least two very different possibilities: the inference is broken, or the inference is fine but one of the premises is false. These two bugs require different fixes.

For example:

```
P1: People with experience are usually more reliable than people without it.
P2: Dorian has less experience in this matter.
C : Therefore Dorian's view is unreliable.
```

Even if we temporarily accept P1, the conclusion still moves too quickly, because "usually more reliable" does not mean "more reliable on every specific proposition." And if we inspect P1 itself, it also needs conditions: what kind of experience? Is it relevant to the proposition under discussion? Is the experience systematically biased? Is there a more direct measurement that can override it? The real issue cannot be resolved by the slogan "respect experience" or "do not worship experience." The premise has to be unpacked.

More troublesome still, many premises are never stated at all. They sit underneath the sentence and run as defaults. Suppose someone says:

> "Young people should listen to those who have been through more."

That statement may secretly depend on an entire hidden stack:

```
older -> more experience
more experience -> better judgement
better judgement in one domain -> better judgement in this domain
past environment ~= current environment
social convention -> normative obligation
```

If even one of these links fails, the conclusion may need to be recompiled. 💀

This is why some conversations produce such a strange form of exhaustion. You think you are refuting one sentence, but in reality you have touched an entire hidden dependency tree. You resolve the first layer, and the other person generates the same conclusion again from the next layer, like a recursive include that never finishes.

So when a discussion remains stuck for a long time, I now start asking:

> **Are we actually disagreeing about the premises rather than the conclusion?**

That question is extraordinarily useful, because it can turn statements such as "you just don't understand," "you're too idealistic," or "you lack experience" back into inspectable form. For example:

"You're too idealistic" may actually mean:

> I think participants in this institution will pursue their own interests more aggressively than you assume, and you are giving that factor too little weight.

Now we finally have something to discuss. At least we know where the disagreement lives.

Premises also have another very important category: **value premises**. Some debates look factual on the surface but are really disputes about what we consider more important. For example:

> "Is this policy a good policy?"

That is not a question from which a set of statistics can automatically output an answer, because you still have to decide:

```
Do we value efficiency more, or fairness?
Do we care more about average outcomes, or the condition of the worst-off?
Do we value short-term stability more, or long-term freedom?
How much risk are we willing to accept?
```

Two people may share exactly the same empirical data and still reach different conclusions because their value premises differ. In that case, bombarding the other person with more data may not solve anything, because the disagreement is not on the data layer.

So I increasingly care about this order:

```
Find the premises first
Then inspect the inference
```

```
Then inspect the evidence
Only then inspect the conclusion
```

Rather than seeing a different conclusion and immediately trying to "prove the other person wrong."

This order also prevents a common illusion: **if the form of the reasoning is elegant, the conclusion must be reliable.** Formal logic can only guarantee that, under true premises, certain forms of reasoning preserve truth. It does not guarantee that the premises themselves are true. `valid` and `sound` are not the same thing. This matters enormously in ordinary conversation.

In engineering language:

```
clean inference over garbage premises
= garbage with excellent architecture
```


And premise-checking has to be applied to myself as well. Often I look at someone else's conclusion and wonder, "How could anyone possibly arrive there?" only to discover that I have also preloaded several assumptions I never stated explicitly. I simply failed to notice them because they matched my own worldview. The real problem is not having premises; every chain of reasoning needs them. The problem is **treating your own premises as the world's default configuration**.

So now I try to express some important judgments conditionally rather than absolutely:

> If assumptions A, B, and C hold, then for now I lean toward X.

This sounds less dramatic than "X is simply true," but it provides one extremely important interface: if A, B, or C changes, I know where the update should occur. I do not need to wait until the entire conclusion collapses before attempting emergency repairs.

In a sense, this is very close to a mathematical theorem. Mathematics does not merely throw "therefore X" at you. It explicitly tells you under which assumptions X holds. Clean thinking should have the same habit:

> **Do not only ask, "How did you derive this conclusion?" Ask also, "Which premises are you allowing yourself to start from?"**

So before entering a long discussion, there is one question I now consider especially valuable:

> **What evidence would make the other person change their mind?**

If the answer is "no evidence at all," then this is not an epistemic debate. It may be identity defense, status negotiation, value signaling, or simply "I cannot admit I was wrong in front of you." In that case, adding more evidence often just increases the data-transfer rate while the actual bottleneck remains in the protocol.

7. What should we do with questions that cannot be verified or falsified?

There is a more difficult boundary here. Some questions are genuinely hard to confirm or falsify through direct experiment. If my only tool is "is there an experiment?", then yes, I become rather powerless.

But I increasingly think this does not mean reasoning has failed. It means we need a different toolset. Empirical claims can be tested through observation and experiment; conceptual claims require definitions, distinctions, and logical consistency; normative claims require explicit value premises followed by valid reasoning; and for some metaphysical or currently undecidable claims, the most reasonable state is often not choosing a side quickly, but preserving uncertainty.

For example:

> "Everything that happens is deliberately arranged by the universe."

If every possible outcome can be reinterpreted afterward as "part of the plan," then the claim provides no observational condition that would let me distinguish between "it is true" and "it is false." I do not need to prove it wrong. I only need to say:

> **The current evidence is insufficient to distinguish whether this claim is true or false.**

Its epistemic status may be:

```
underdetermined
```

rather than:

```
false
```

That restraint matters. Scientific training can sometimes create the illusion that every question should eventually output `true` or `false`. But quite often, the most precise answer is:

```
insufficient information
```

That is not failure. It is refusing to pretend to know.

And "unfalsifiable" should not be simplified into the magic formula "therefore false." A more precise statement is this: if a claim does not specify what observation would count against it, then it is difficult to treat the claim as an ordinary scientific proposition with empirical content. It may still function as a metaphysical belief, a literary worldview, a religious claim, or an attitude toward life. But at that point, the rules of discussion have changed. The deepest confusion usually comes not from the coexistence of different kinds of claims, but from moving the language of one category into another while refusing to admit that the protocol has switched.

For example, "I choose to believe that life has some kind of meaning" and "the universe contains an objectively verifiable mechanism that arranges events according to the frequency of human thought" are entirely different claims. The first primarily expresses a stance or life framework. The second begins making causal claims about the external world. Once you move into the second category, you cannot enjoy the authority conferred by scientific-sounding language and then retreat to "you just haven't awakened enough to understand" the moment verification is requested. That two-way escape pod is extremely convenient, which is precisely why it deserves suspicion.

8. "Other people do not understand me" and "I actually explained myself clearly" are not the same thing

Looking back at my old feeling that "nobody understands me," I now think part of it was real and part of it was simply unfinished communication.

An idea has not truly been explained merely because I feel I have spoken it. It has been explained when another person can reconstruct a sufficiently similar structure in their own mind. That means what I need to preserve is not the original sentence, not my favorite metaphor, and not some elegant technical term. It is the invariant.

Take the "cache miss" analogy. The actual structure I may want to preserve is simply:

> The first attempt is expensive; once a model exists, the cost goes down; differences in cost do not necessarily imply differences in ability.

That structure can be expressed through computing, walking, learning an instrument, finding a file, or cooking. Advanced communication is not about discovering the one uniquely correct sentence. It is about **finding, for different receivers, the encoding most likely to let them reconstruct the same structure.**

From this perspective, communication resembles a cross-model translation problem. You cannot assume both sides share the same vocabulary, priors, background knowledge, or even the same standards of judgment. Difficult communication often runs all of these at once:

```
content generation
+ audience modelling
```

```
+ error detection
+ translation
+ epistemic negotiation
```

Five threads running simultaneously. Of course the fans take off. 💀

This also changed how I understand the phrase "being good at explaining things." It is not merely having a large vocabulary or writing beautiful sentences, and it is not even merely being logically correct. The difficult part is estimating which dependencies the other person already has, which ones are missing, and then supplying only what is necessary. Say too little and they cannot decode it. Say too much and working memory goes OOM. A good explanation is not a dump of everything I know. It is an attempt to let another person reconstruct a similar model in their own cognitive system while adding as little unnecessary load as possible.

From this perspective, analogy is no longer just decoration. It is a temporary adapter: it maps an unfamiliar structure onto something the other person already has. But adapters also introduce abstraction loss. `cache miss` may be excellent for expressing "an existing model reduces future processing cost," but it does not imply that the human brain literally works like a CPU cache. If I forget that, a useful analogy quietly begins mutating into a theory. That is exactly the same problem I criticized earlier when talking about "frequency" and "energy," so the blade has to cut in my direction too. 😇

9. Not every conversation is worth running to convergence

I used to assume more readily that, once a conversation had begun, we ought to make everything clear, and ideally both sides should eventually reach some kind of agreement. I increasingly think that goal is too expensive, and often unnecessary.

Some conversations are simply about getting along and do not require epistemic convergence. Some are exchanges of viewpoints, where clarifying two or three layers is enough. Some involve shared decisions, in which case aligning models is genuinely worthwhile. Some people just want to express themselves and have never invited you to correct them. Treating every interaction as a session of "jointly converging toward truth" is neither realistic nor computationally cheap.

More importantly, I now try to determine first: **is this actually a truth-seeking context?**

If it is, then keep debugging: inspect definitions, evidence, counterexamples, and inference. If it is not, then failing to persuade the other person does not have to be interpreted as "my model failed." At the same time, another firewall is needed: the other person's refusal to update does not mean I am automatically correct. Two people can absolutely consist of one stubborn person and one person who is wrong with exquisite sophistication. 🤣

So I now prefer a more stable stance:

I allow the other person not to update, and I allow myself to be wrong.

There can also be a stop condition. If a conversation repeatedly involves redefining the other person's position, refusing to specify any condition under which one's own view could change, absorbing every counterexample into the theory after the fact, or shifting from arguments to identity and credentials, then continuing is usually no longer a high-yield truth-seeking activity. Ending the discussion does not mean agreeing or admitting defeat. It simply means I am no longer donating unlimited CPU time to an iterator with no convergence condition.

Conversely, if someone can say, "If that evidence holds, I would reduce my confidence from 80% to 50%," then even if we still disagree in the end, I would consider it a high-quality discussion. Perhaps what matters is not synchronizing to the same answer, but ensuring that both models retain an update path.

I want my mistakes, at least, to be locatable, modifiable, and capable of being downgraded in confidence, rather than untouchable conclusions maintained by pride, identity, or beautiful language.

If I had to compress this entire experience into one sentence, it would probably be:

> **I did not stop liking abstraction; I started demanding that abstraction have a schema. I did not stop trusting intuition; I started requiring intuition to carry a confidence level. And I stopped assuming I had to persuade everyone before first asking whether the conversation was worth the communication cost.**

In the past, a sentence like "Understand the energy of the universe and you will understand life" might have held my attention for a while. Now my first reaction is closer to:

```
type error: "energy" undefined
```

followed by:

```
please provide definition, mechanism, evidence, or at least a decent example
```

If none of those arrive, then the claim can go back on the shelf for now. 📚😈

10. What I really want to preserve may not be "rationality," but updatability

By the time I reached the end of this piece, I realized that what I truly value may not be an identity of "I want to be more rational than other people." That identity is dangerous in its own way. Once "I am a rational person" becomes a self-image that needs protection, I can start defending my conclusions too. I may simply replace "Professor So-and-so said it" with "I researched this," "I understand statistics," or "I have read Spivak." The presentation becomes more elegant, but the bug can remain exactly the same.

What I want to preserve more is **updatability**. A view can be strong. A model can temporarily look excellent. But ideally it always keeps a few interfaces open: what evidence would make me revise it? Which parts are supported by data, and which parts are my extrapolation? Which concepts are merely useful abstractions rather than the world itself? If I change the assumptions, does the conclusion change with them? Am I describing facts, or expressing values?

Perhaps this is where mathematics, writing, and physics eventually converge for me. On the surface, they teach completely different things: one deals with proof, one with language, one with natural phenomena. Yet all three keep reminding me of the same principle: **do not disguise an unexpanded intuition as a complete model, do not disguise elegant compression as proof, and do not exempt an idea from inspection simply because it belongs to you.**

That may be the thing that made it impossible for me to go back. When I see a claim now, I find it difficult to ask only, "Does this sound reasonable?" A few extra checks start running automatically:

```
What exactly is the claim?
What premises does it depend on?
Which premises are empirical, and which are value judgements?
What do the words mean here?
Does the conclusion actually follow from the premises?
What evidence changes its probability?
What would count against it?
Where did abstraction lose information?
Am I using a metaphor, or claiming a mechanism?
```

These questions sometimes make life slightly more inconvenient. Quite a few pleasant pieces of "life wisdom" now produce an entire screen of warnings during local compilation. 🪆 But in exchange, there is another kind of freedom: I do not have to classify everything immediately as true or false, and I do not have to rent someone else's certainty just because they speak with confidence. I can say, "I don't know," "I only believe this part with 60% confidence," "this is a good analogy, but do not push it into a mechanism," or "we may simply not have enough information yet."

If this counts as a kind of "protagonist-style thinking," then it is not the belief that the universe will arrange the plot for me. It is accepting something far more ordinary, and far more difficult:

Your model is yours to maintain. 🧠🔧