Standing in the Middle

By William Tygart · November 13, 2026

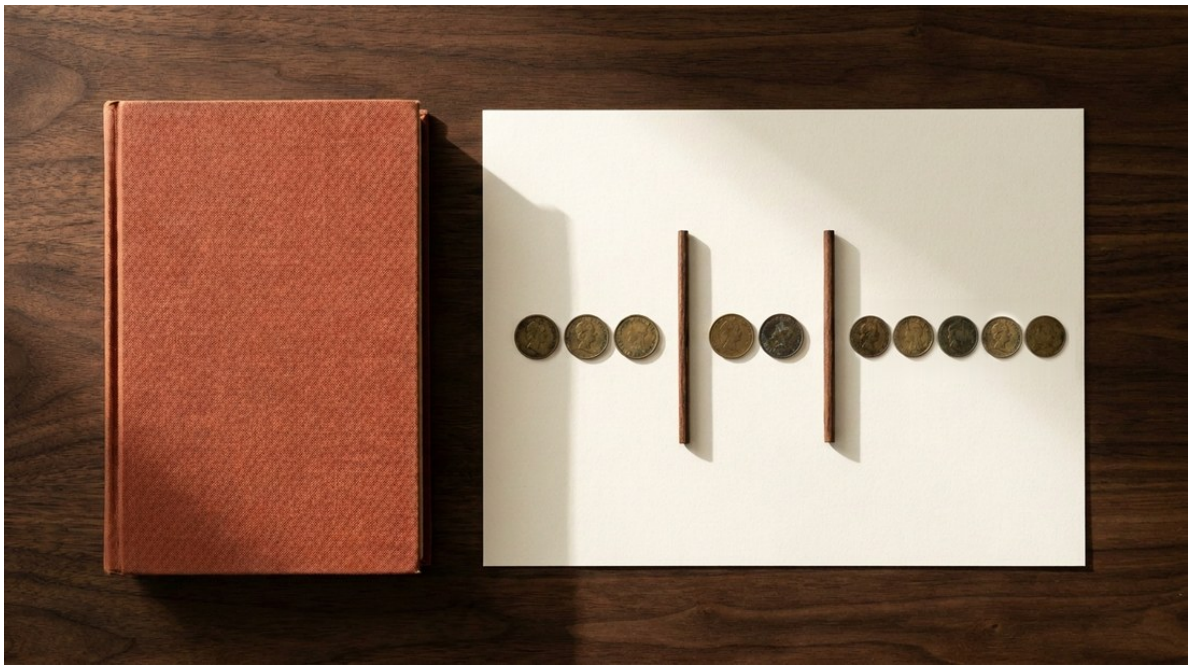
One direction gives you something you can point to. The other gives you a wider map. Meaning shows up only where the two have to agree.

This morning I sat down with Chapter 2 of Douglas Hofstadter's *Gödel, Escher, Bach: An Eternal Golden Braid* (*"Meaning and Form in Mathematics"*), right after finishing the `MU`-puzzle and circling back to the short prelude adapted from Lewis Carroll's *What the Tortoise Said to Achilles*.

In Carroll's dialogue, Achilles keeps trying to force the Tortoise to accept a conclusion by writing down one more rule as another line in his notebook — and the Tortoise happily accepts every rule and still refuses to budge, trapping them in an endless chain. You cannot push meaning forward by piling more marks inside the same closed system. Something outside has to line up with it.

Then Hofstadter hands you a tiny formal toy called the `pq`-system. Strings built out of three characters only: `p`, `q`, and hyphens, like `---p--q-----`.

My first reaction was to ask whether the hyphens were just noise I needed to clear away to find the signal. They aren't noise. They are the count: three hyphens on the left, two in the middle, five on the right. Put ten tarnished brass coins on a sheet of paper with two sticks between them — three coins, a stick, two coins, a stick, five coins — and you see `3 + 2 = 5`.



Three coins, a stick, two coins, a stick, five coins: the marks stay mute until a structure lines up with them.

Here is the catch that stopped me: the coins don't know they are doing arithmetic, and the hyphens don't know they mean three, two, and five. Meaning isn't packed inside the symbols. It lives in the fit — the

isomorphism — between the marks sitting on the table and an outside structure that lines up with them.

That single idea cracked open a frustration I have been carrying for months.

Why I used to roll my eyes at keynote stages

I work bottom-up. My week happens inside origin server logs, backup checksums, and request counts — the actual exhaust left behind when systems run. Everything we put into our white paper, *Not All Bots Are the Same: 55 Days of Server Logs from One Website*, and our log series like *My Analytics Said 1,000 Visitors. My Logs Said Otherwise.* came from counting raw requests on the server instead of trusting a dashboard estimate.

Because of that, watching SEO speakers on keynote stages used to drive me crazy. From where I sit, a lot of stage talk sounds like pure top-down storytelling: wide, inferred narratives with nothing you can trace back to a timestamp or a status code. Bottom-up has actuals. Top-down has inference.

Then, halfway through talking out `---p--q-----` this morning, I caught my own framing.

Looking at that string, I realized I might just be the three hyphens on the left — sharp on my own local count, blind to the rest of the equation. The people on stage aren't living inside one log file, which is precisely why they can hold a wider landscape across variables I cannot see from down inside the requests.

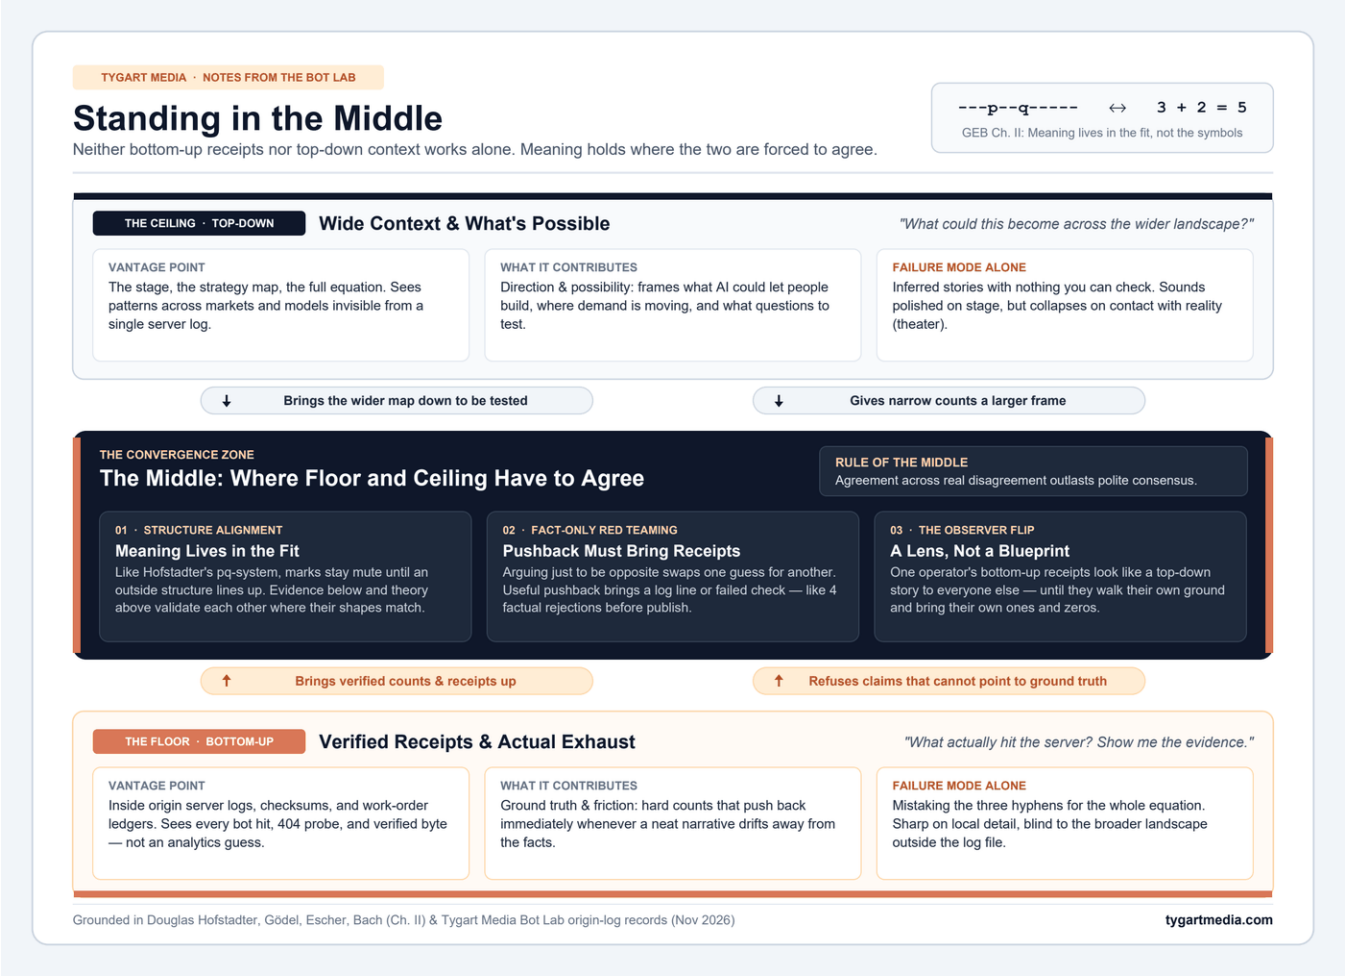
Neither direction is superior. There is no "better." Bottom-up gives you something you can point to. Top-down gives you a wider map. If you already have the bottom, you don't need more bottom — you need to look up. If you only have the top, you don't need another keynote story — you need to climb down into the ground truth.

The floor, the ceiling, and the middle where they meet

I have been calling the two ends **the floor** and **the ceiling**.

- **The floor** is what actually happened: the `2,838,206` requests in 55 days of `tygartmedia.com` server logs (`Sep 18–Nov 12, 2026`), the HTTP status code, the work-order accounting rule where unfinished work stays at `$0` until verified (*Just a Guy: What Is One Person Worth Now?*).
- **The ceiling** is what is possible: what AI can do, where retrieval and buyer behavior are heading, the broader theory you want to test next.

Most real thinking doesn't live on the floor or up against the ceiling. It lives **in the middle**, where the two are forced to agree.



The Floor (bottom-up receipts), The Ceiling (top-down context), and The Middle where the two are forced to agree.

When the floor and the ceiling meet in the middle, neither side gets a free pass. A wide narrative has to point to something that actually hit the server. A raw log count has to explain what it means outside its own three hyphens. Just like Hofstadter's pq-system, meaning shows up only where the two structures lock together.

Agreement with receipts beats theater every time

Once you see that, something else clicks into place: an agreement reached across real disagreement is stronger than polite agreement you started with.

That is why the U.S. Constitution splits power across three branches that push back against each other. Picture two axes crossing at an angle: they can diverge everywhere else, and right where they intersect is the point that carries the load.

It is also the real difference between useful red teaming and performative pushback. I have always disliked red teaming when it is just someone taking the opposite side for the sake of being opposite. That is inference fighting inference — pure theater — and if you listen closely to critics, empty pushback can talk you right out of a sound idea.

Fact-only pushback is the opposite. When we put together the PDF package for *Not All Bots Are the Same*, it went back two, three, four times over on verification checks before it was ready to publish. In the middle of

those rounds, it felt annoying and disappointing. Looking back afterward, it was the best part of the process: every rejection came back with *"show me the evidence."* The bottom-up receipts forced the top-down story to earn its keep.

The catch: my bottom-up is your top-down

Here is the turn I didn't see clearly until the end of the morning: **what is bottom-up for me is still top-down for you until you walk it yourself.**

Even when I publish the exact counts from our logs or lay out how a week of work orders ran, to anyone reading this article on a screen, my receipts arrive as someone else's narrative. My Tuesday is not even the same as my Thursday, because things change — and your ones and zeros, your site, your crew, and your moment are specific to you.

That is why I am not handing you a step-by-step checklist, and I am not selling a method. Think of this as a report from where I am standing:

- **If you already live up at the ceiling** — in strategy decks, big-picture theses, and wide industry maps — maybe this gives you a reason to climb *down* into the raw receipts and see what pushes back.
- **If you already live down on the floor** — buried in logs, tickets, and daily counts — maybe this gives you permission to look *up* at the wider landscape without feeling like you are abandoning the facts.

I am just reporting from the field. Put your own spice on it.

We do AI citation restoration. When the AI can't find you, can't read you, or keeps citing somebody else, send us the worst one you have — your RFPs, your disasters, the problem everyone else already failed to fix. We come from the restoration industry. We know what's on the floor, and cleaning it up is the job. And if you're a developer who just needs a hand on the hard part, get a hold of us. That's what we're here for.

— William Tygart, Tygart Media. Reflections from a morning reading session on Gödel, Escher, Bach (Ch. II) and Bot Lab origin-log records, November 13, 2026.