

Doubting the Machine

*An honest reckoning with AI — for people
who aren't sold on the hype, or the doom.*

SKEPTIC'S EDITION

Dr. Ayse Ozturk

Clinical Associate Professor of Marketing · Darla Moore School of Business, University of South
Carolina

drayseozturk.org

Your doubt is not a deficiency. It's the most useful thing here.

Most writing about AI is trying to sell you something — a product, a worldview, a wave of inevitability you're supposed to ride or be left behind by. The boosters tell you it changes everything; the doomers tell you it ends everything; both are loud, and both have something to gain from your certainty. If you've watched all this with arms crossed, you have been the sane one in the room.

This edition is written for that person, and it isn't here to talk you out of your skepticism. It's here to aim it. There are real reasons for concern — about accuracy, jobs, the environment, bias, copyright, and your privacy — and there's also a lot of nonsense in both directions. The useful skill is telling them apart, and that requires exactly what you already have: a refusal to take claims on faith. Each chapter takes one of the big worries, states it as strongly as its believers would, and then walks through the actual evidence — validating what holds up and puncturing what doesn't.

I'm an academic, not a salesperson, and I've linked a source for every factual claim so you can check my work rather than trust my tone — which is, after all, the whole point. Everything draws on my notes at drayseozturk.org. Stay skeptical. Just be skeptical of the hype *and* the panic.

Ayse

Dr. Ayse Ozturk · Columbia, South Carolina

Five honest questions

Each worry, stated fairly — then weighed against the evidence.

01 **Is it just hype?**

It makes things up and the demos oversell — both true. And it measurably helps real work. The honest answer is "neither magic nor nothing."

02 **Will it take my job?**

The disruption is real and the "it'll replace everyone" story is overblown. What the labor evidence actually shows about who's exposed.

03 **Is it wrecking the planet?**

One question costs almost nothing; a billion questions plus the buildout costs a great deal. Where the real environmental concern lives.

04 **Is it biased, built on stolen work, and spying on me?**

Three serious objections — bias, copyright, privacy. Which are proven, which are unresolved in court, and what you can actually do.

05 **So should a thoughtful person just opt out?**

Refusing has costs too. The case for engaged skepticism over both the cheerleading and the boycott — and why your doubt belongs at the table.

01

QUESTION ONE

Is it just hype?

The boosters say it changes everything; the critics say it's a glorified autocomplete that lies. Strip away both and a more useful, more boring truth remains.

Let's start with the strongest version of the skeptic's case, because it has real teeth. These systems don't understand anything; they predict likely text. They state falsehoods with total confidence — and not rarely. The product demos are staged, the founders have billions riding on your belief, and the gap between the keynote and your actual Tuesday is wide. Every bit of that is true, and anyone who waves it away is not being honest with you.

The fabrication problem is worse than the boosters admit. On its own factual-accuracy test, OpenAI found its newer, more advanced models hallucinated *more* than the older ones — a third to nearly half the time. "It'll fix itself in the next version" is not supported by the evidence; sometimes the next version is worse on exactly this.

Skepticism about the hype is correct. Dismissing the usefulness isn't. Both are true at once.

And yet. When researchers actually measure whether it helps real people do real work — not in a demo, in controlled studies — the effect is consistently, unglamorously real. Professionals doing everyday writing finished about 40% faster, with quality going up. Customer-support workers handled more cases, with the biggest gains going to the least experienced. You can hold both facts in one hand: it is genuinely flawed *and* genuinely useful. The hype is overblown; the dismissal is also wrong. The honest position is the boring middle, and it's the only one the evidence supports.

**up to
48%**

How often OpenAI's *newer* models fabricated answers on its own factual test — higher than the older model. "It'll fix itself" isn't supported.

OpenAI system card, 2025

40%

faster on real writing tasks in a controlled study — quality up, not down. The usefulness is measurable, not just marketing.

Noy & Zhang, Science, 2023

You're not the only doubter

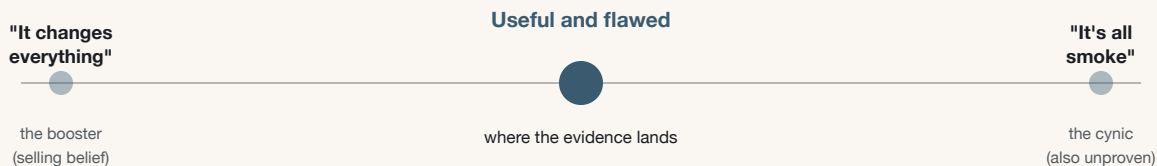
If your skepticism has felt lonely, the data says otherwise. Stanford's annual AI Index — about as neutral a scorekeeper as exists — reports that as adoption climbed, recorded AI-related incidents hit a record high and public trust in AI companies *declined*. Doubt isn't a fringe position held by people who "don't get it." It's a widespread, evidence-backed response to a technology that has been oversold and under-governed.

So the move isn't to pick a tribe. It's to downgrade every confident claim — in either direction — until you've seen the evidence. Treat the booster's "it changes everything" and the cynic's "it's all smoke" as equally unproven until shown otherwise. That habit, applied consistently, is what the rest of this booklet is.

Trust is falling, not rising.

As use grew, public trust in AI firms dropped and reported incidents hit a record. Your doubt is the mainstream, evidence-backed position.

Stanford HAI AI Index, 2025



The honest position is the least exciting one — which is exactly why no one's selling it.

NEITHER POLE SURVIVES CONTACT WITH THE EVIDENCE. It fabricates and the hype is real; it also measurably helps. Skepticism's job isn't to pick the opposite extreme — it's to find the calibrated middle and stay there.

HOW TO WEIGH IT

- 1 Discount the source's stake.** When a claim comes from someone who profits from your belief — or your fear — raise your bar for evidence accordingly.
- 2 Separate "impressive" from "reliable."** A dazzling demo says nothing about whether it works on your task, every time. Ask for the second, not the first.
- 3 Trust measured results over narratives.** Controlled studies beat keynotes and doom-threads alike. When neither exists yet, withhold judgment.

QUESTION TWO

02

Will it take my job?

"It'll replace everyone" and "it'll change nothing" are both wrong. The labor evidence points somewhere more specific — and more manageable.

This is the fear that keeps people up at night, and it deserves to be taken literally, not soothed away. The strong version: these tools are getting better fast, they do the writing-and-analysis tasks that white-collar jobs are made of, and companies have every incentive to use them to employ fewer people. Some of that is already happening. Anyone who tells you "AI only creates jobs, never destroys them" is reciting a slogan, not reading the data.

But the data, when you read it, tells a more specific story than "it replaces everyone." The World Economic Forum's 2025 survey of employers projects roughly 170 million new roles created and 92 million displaced by 2030 — enormous churn, a net increase, and a great deal of disruption for the people on the wrong side of that churn. The same report finds AI is the fastest-growing skill employers want, and — in the same breath — that analytical thinking remains the single most-valued human skill. The pattern isn't mass replacement. It's reshuffling, fast and unevenly.

The honest fear isn't replacement. It's **churn** — and being on the wrong side of it.

What actually protects you tracks that pattern. In study after study, AI doesn't replace the worker who can judge, decide, and verify — it replaces the narrow task. The people most exposed are those whose role is only the task AI now does; the people who gain are those who direct the tool and own the judgment around it. That's not a comfort, exactly, but it is a direction.

+78M

Projected net new jobs by 2030 — about 170M created against 92M displaced. A net gain, and also enormous churn for those displaced.

WEF Future of Jobs, 2025

Most valued skill: still human. The same employers naming AI the fastest-growing skill name analytical thinking the most essential one. Judgment isn't the thing being automated.


The honest, uncomfortable nuance

Two things are true that the slogans can't hold together. The macro picture is reshuffling, not apocalypse — historically, automation has moved labor rather than ended it, and the current evidence fits that shape. *And* the macro picture is cold comfort if you're one of the 92 million, in a specific job, in a specific year. "Net positive over the decade" does not pay a mortgage during the transition. A serious skeptic holds both: refuse the panic, and refuse the breezy reassurance that waves away real, concentrated harm.

The defensible response, then, is neither to ignore the threat nor to despair at it. It's to make yourself the person who directs the tool rather than the person whose whole job it now does — and to support the policies and retraining that catch the people churn lands on. Skepticism here means demanding that "net positive" come with a plan for the displaced, not just a press release.

Both, at once. "Net jobs gained" can be true while real people lose specific jobs. Reject the slogan in either direction.

Projected global jobs by 2030 (employer survey)



A net gain — and 92 million people on the losing side of the churn. Both numbers are real.

CHURN, NOT APOCALYPSE — BUT CHURN IS REAL. The net is positive; the disruption underneath it is enormous and unevenly distributed. The honest reading holds the reassuring number and the painful one together. Source: WEF Future of Jobs, 2025.

HOW TO WEIGH IT

- 1 Ask "task or judgment?"** The roles most exposed are pure task. The most protected add judgment, context, and accountability the tool can't supply.
- 2 Refuse both slogans.** "Replaces everyone" and "creates only" are each marketing. The truth is large, uneven churn — plan for it, don't deny it.
- 3 Demand the plan, not just the projection.** Hold companies and policymakers to retraining and support for the displaced — that's where skeptical pressure does real good.

QUESTION THREE

03

Is it wrecking the planet?

The viral claims about one query draining a bottle of water are mostly wrong. The real concern is bigger, slower, and harder to dismiss.

This is where the skeptic has to be skeptical in both directions, because the environmental conversation is full of bad numbers — and a real problem underneath them. Start with the viral version: that a single AI question guzzles a bottle of water or a phone charge of power. On the best figures we have, that's simply not right. One query uses roughly the electricity an oven draws in about a second, and a tiny fraction of a teaspoon of water. As a per-use number, it's close to noise.

But — and this is the part the industry would rather you stop reading — those are the companies' own figures, not independently peer-reviewed, and they exclude the enormous energy used to *train* the models in the first place. And small times gigantic is not small. Multiply a near-nothing by a billion queries a day, then add a global construction boom in data centers, and you get a genuine problem.

The problem was never your one question. It's a billion of them — plus the buildout behind them.

The neutral scorekeepers confirm the scale. The International Energy Agency projects electricity use from data centers will roughly double by 2030, to around 3% of all global electricity, with AI the fastest-growing slice. That's the honest environmental story: not your individual chatbot use, which is trivial, but the aggregate demand and the infrastructure behind it, which is large, real, and growing fast.

~0.34_{wh}

Energy per query, by the maker's own estimate — about a second of an oven. Per use, it's near noise. (Their figure, not peer-reviewed; excludes training.)

via Data Center Dynamics, 2025

~2×

Projected rise in data-center electricity use by 2030 — to ~3% of global demand, with AI the fastest-growing part. The aggregate is the real issue.


IEA, Energy and AI, 2025

Aim the concern where it counts

This matters because misdirected guilt is wasted leverage. Agonizing over whether to ask a chatbot one more question is, environmentally, like refusing a single napkin to save a forest — the per-use cost is genuinely tiny, and the hand-wringing changes nothing. The decisions that actually move the needle are upstream and collective: how data centers are powered, whether they're built where water is scarce, whether companies are honest about their full footprint including training, and whether the grid they draw on is being decarbonized.

So the skeptical environmental position isn't "never use it" or "it's fine, relax." It's to drop the per-query guilt, and aim your scrutiny at the companies and policies that control the aggregate — demanding transparency on the numbers they currently keep vague, and clean power for the buildout that's coming either way.

Where your leverage is. Not your query count — that's noise. It's pressure for clean-powered data centers and honest, full-footprint reporting.



TWO VERY DIFFERENT SCALES. Per use, the cost is trivial and your guilt is misplaced. In aggregate — a billion queries a day plus a building boom — it's a real and growing problem. Aim your concern at the scale that matters. Sources: per-query, company estimate; aggregate, IEA 2025.

HOW TO WEIGH IT

- 1 Drop the per-query guilt.** Your individual use is environmentally trivial; agonizing over it is wasted energy of a different kind.
- 2 Distrust both the viral stat and the company brochure.** "A bottle per query" is wrong; "totally green" is unproven. Ask for independent, full-footprint numbers.
- 3 Push upstream.** The leverage is in how data centers are powered and sited — that's where a skeptic's pressure actually counts.

QUESTION FOUR

04

Is it biased, built on stolen work, and spying on me?

Three of the strongest objections, and the honest answers differ: one is proven, one is unresolved in court, one is partly in your hands.

These three get lumped together as "AI is unethical," but they have genuinely different evidence behind them, and a skeptic should grade them separately rather than accept or dismiss the bundle.

Bias is real and measurable. These models learn from human text and absorb its slants; this isn't speculation, it's documented — including by the makers themselves, who now publish tests scoring how evenhandedly their models treat contested topics. The honest takeaway isn't "AI is hopelessly biased" or "it's perfectly neutral." It's that bias is a measured, managed quantity that varies by model and topic, and you should treat any single answer as a perspective, not the truth.

The copyright objection is serious and unresolved. These systems were trained on vast amounts of text and images, much of it copyrighted, much of it without the creators' consent or payment. Whether that's legal is being fought out right now — dozens of lawsuits, including The New York Times against OpenAI, with judges so far split on whether it counts as fair use and no settled answer expected before 2026. Anyone who tells you the training question is resolved, in either direction, is ahead of the courts.

The copyright question is genuinely unsettled. Anyone who tells you otherwise is ahead of the courts.

Privacy is partly in your hands. Free tools may use what you type to improve their systems, and your conversations aren't guaranteed to be private. That's a real concern — and also the one you have the most direct control over, by what you choose to type.

Bias: proven. Makers' own evaluations score how evenly models treat contested topics — confirming bias exists and varies. Read any answer as a view, not the verdict.

Anthropic, 2025

Copyright: in court. Dozens of suits (incl. NYT v. OpenAI) remain undecided; judges are split on fair use. Genuinely unresolved.

AI litigation tracker, 2025

Grade each concern, then act on it

The reason to separate these is that they call for different responses. Bias you manage by reading critically — cross-check contested claims, ask for the opposing view, never treat one model's answer as neutral fact. Copyright you can't personally resolve, but you can take seriously: prefer tools and companies that license their data or compensate creators, and support the creators caught in the middle. Privacy you control directly: keep anything sensitive — personal, medical, financial, confidential — out of the chat box, and use the privacy settings and opt-outs the tools provide.

None of this requires either swallowing the technology whole or rejecting it outright. It requires doing what skeptics do best: refusing the convenient bundle, weighing each claim on its own evidence, and acting in proportion to what that evidence actually supports.

Privacy: your move. Treat the chat like a postcard. The control you have here is real — use it, and use the opt-outs.

Bias

PROVEN

Real, measured, varies by model.

Your move:
read critically,
cross-check.

Copyright

UNRESOLVED

In court; judges split; undecided.

Your move:
favor licensed tools; back creators.

Privacy

IN YOUR HANDS

Inputs may be used; not guaranteed private.

Your move:
keep sensitive data out; use opt-outs.

THREE OBJECTIONS, THREE HONEST VERDICTS. They don't deserve the same answer. Bias is proven and managed by reading critically; copyright is genuinely unsettled; privacy is largely within your control. Grade them separately. Sources: Anthropic 2025; ongoing litigation.

HOW TO WEIGH IT

- 1 Unbundle "it's unethical."** Grade bias, copyright, and privacy on their own evidence — they're at very different stages of proof.
- 2 Read every answer as a perspective.** For anything contested, ask for the opposing case and cross-check before you trust it.
- 3 Use the control you do have.** Keep sensitive information out of the chat box, turn on the privacy opt-outs, and favor tools that treat creators fairly.

QUESTION FIVE

05

So should a thoughtful person just opt out?

Refusing has a logic — and a cost. The case for engaged skepticism over both the cheerleading and the boycott.

After all that — the fabrication, the churn, the energy, the unresolved ethics — the honest skeptic reaches a fair question: why engage at all? Why not simply refuse, on principle, and wait for the dust to settle? It's a respectable position, and it's not the same as ignorance. Plenty of thoughtful people are choosing distance, and they have reasons this booklet has just spent four chapters taking seriously.

But refusal has costs of its own, and a real skeptic counts those too. A technology this consequential is going to be shaped by someone — in workplaces, schools, courts, and policy — and the people who opt out of understanding it also opt out of the conversation about how it's governed. Worse, blanket refusal cedes the ground to exactly the two groups you trust least: the boosters selling it and the companies building it. Your informed doubt is more valuable inside the room than outside it.

Opt out of understanding it, and you opt out of the **vote** on how it's used.

There's a middle path between the cheerleader and the abstainer, and it's the one this whole edition has been modeling: engaged skepticism. Understand the tool well enough to judge it. Use it where the evidence says it genuinely helps, with eyes open to its flaws. Refuse it where the costs outweigh the benefits. And keep demanding better — more honesty, more accountability, better governance — from a position of knowing exactly what you're talking about.

Refusal isn't neutral. Sitting out doesn't stop the technology — it just removes informed critics from the decisions about it.

The most valuable seat. An informed skeptic shapes how AI gets used far more than either a cheerleader or an absentee.

Skepticism is a method, not a verdict

Here's the reframe worth leaving with. Skepticism was never supposed to be a conclusion — "I'm against it" — fixed in advance. It's a method: withhold belief until the evidence earns it, then believe exactly as much as the evidence supports, and no more. Applied to AI, that method doesn't deliver a verdict of "good" or "bad." It delivers a posture — calibrated, alert, hard to fool in either direction — which is precisely what this moment is short of.

So keep your doubt; it's the most valuable thing you brought. Just point it at everyone equally: the founder's keynote, the doomer's thread, the company's sustainability page, and this booklet too. Check the sources. Weigh the evidence. Change your mind when it changes, and not before. That's not opting out. That's the most engaged, most useful thing a thoughtful person can do with a technology this loud.

Doubt everyone equally.

The boosters, the doomers, the companies — and this booklet. Check the sources; that's the method working.

Evangelist

Adopts everything,
believes the demo.

Misses the flaws.
fooled by hype

Engaged skeptic

Understands it,
uses what works,
refuses what doesn't,
demands better.

**hard to fool —
and in the room**

Absentee

Refuses on principle,
opts out entirely.

Loses the vote.
cedes the ground

THREE STANCES — ONLY ONE IS CALIBRATED. The evangelist is fooled by hype; the absentee cedes the decisions to others. The engaged skeptic understands the tool, uses what's proven, refuses what isn't, and keeps the pressure on. That's where your doubt does the most good.

HOW TO WEIGH IT

- 1 Treat skepticism as a method, not a side.** Withhold belief until the evidence earns it — then believe exactly that much, in either direction.
- 2 Stay in the room.** Understand the technology well enough to shape how it's used at your work, your school, your ballot — rather than ceding that to the sellers.
- 3 Doubt this booklet too.** Follow the links, check the sources, and change your mind only when the evidence does. That's the method working as intended.

Where to take this next

Four pages on drayseozturk.org, chosen for the doubter who wants to check things themselves — each with why it earns your scrutiny.

AI Insights Blog

drayseozturk.org/ai-insights-blog

Longer, sourced discussions of the harder questions — where the claims in this edition get unpacked further, with the working shown rather than asserted.

AI Notes

drayseozturk.org/ai-notes

The running source behind this booklet — the hallucination, bias, and sustainability tracking, kept current as the evidence (and the claims) keep changing.

Frontier Models

drayseozturk.org/frontier_models

The actual capabilities and limits of today's leading models, side by side — so you can judge the claims against the tools yourself instead of taking anyone's word.

AI Dictionary

drayseozturk.org/ai-dictionary

Plain definitions to cut through the jargon that hype hides behind — because a vague word is often where an overstated claim is smuggled in.

SOURCES CITED IN THIS EDITION

OpenAI, [o3 and o4-mini System Card](#) (2025) · Noy & Zhang, [Productivity effects of generative AI](#) (Science, 2023) · Stanford HAI, [2025 AI Index Report](#) · World Economic Forum, [Future of Jobs Report 2025](#) · IEA, [Energy and AI](#) (2025) · Data Center Dynamics, [on per-query energy/water](#) (2025) · Anthropic, [Measuring political bias in Claude](#) (2025) · McKool Smith, [AI litigation tracker](#) (2025). Company self-reported figures (energy/water) are noted as such; model and litigation specifics are point-in-time and will move.

DR. OZTURK EDITIONS

Short, honest field guides to working with AI.

Each edition takes the same body of notes and rebuilds it for one audience — skeptics, beginners, professionals, educators, students, small businesses, everyday life, and advanced practitioners — so you read only what's right for you. Clear, concrete, current, and sourced so you can check the work. Made to be finished in one sitting and argued with the next day.

drayseozturk.org

A FREE AI-LITERACY HUB BY DR. AYSE OZTURK · DARLA MOORE SCHOOL OF BUSINESS