

Who Pays If It's Wrong?

Hundreds of billions of dollars are riding on a particular bet about how AI gets better. The bet may well come off. This is about what happens to everybody else if it does not.



WHERE THIS SITS

There is more language in the world than anyone can attend to. The obvious fix, an AI that reads it all for you and tells you what is true, does not work: the judgement that would take requires having been wrong at your own expense, and nothing a model does costs it anything. So the judging stays with people, and almost nothing we currently inspect can tell whether a person actually judged or simply agreed.

01 Why Does It All Feel Like Too Much?

02 Why Can't You Just Be Told?

03 What Holds a Claim Up?

04 How Would You Know?

05 **Who Pays If It's Wrong?**

06 What Would You Build Instead?

Six short pieces, each drawn from a working paper. Each one stands on its own; this is the fifth. They are ordered so that every piece ends by raising the question the next one is called.

You are probably in it already

If you have a pension, a workplace retirement plan, or money in an index fund, you own a piece of this. Not because you decided to back artificial intelligence, but because the funds most people's savings sit in hold the whole market, and a handful of AI-exposed companies have become a very large part of what the whole market is.

Nobody asked you. There was no form. The exposure arrived through the ordinary machinery of saving sensibly.

This is not about whether the bet is a good one. That is genuinely arguable, and the paper says so repeatedly. The question is who is holding it.

What the bet actually says

Underneath the spending is an argument with four steps. It is rarely written down in one place, because each step is usually used to justify the next rather than stated on its own.

01 SCALING WORKS

Bigger models keep getting better as they get bigger.

02 SO AGI FOLLOWS

Or at the very least, a huge jump in what can be produced.

03 SO THE SPENDING IS JUSTIFIED

Hundreds of billions of dollars, committed up front.

04 SO THE WINNER TAKES MOST

Which rewards spending early, and spending hard.

The paper calls this a reasonable bet in the same breath as calling it a bet, and both halves of that sentence are doing work. Every one of the four steps holds only while the first one does.

Four steps, one dependency

The rule underneath is the diagram. The boxes are only its evidence.

THE BET, IN THE ORDER IT IS USUALLY MADE



every one of these holds only while the first one does

You have seen this shape before

What Holds a Claim Up? finds the housing market of 2006 stacked exactly this way.

2006 · THE HOUSING MARKET



2026 · AI INVESTMENT



both chains stand or fall on their first box

Four people who think the first step is wrong

The whole chain rests on step one. So it matters that four well-credentialled researchers, from four different directions, think step one does not hold.

YANN LECUN

A Turing Award winner and for a decade Meta's chief AI scientist. Language models are useful, but are a turning off the road to human-level intelligence.

GARY MARCUS

The reasoning gap is not closing at the rate the scaling story requires, and what looks like reasoning is often retrieval wearing its clothes.

RODNEY BROOKS

When people watch a system do one thing well they assume it can do the neighbouring things, and are usually far too optimistic about that.

FRANÇOIS CHOLLET

Memorisation is useful and is not intelligence. On his benchmark, built to test that, one leading model scored about 5 per cent where people score 84.

And the honest part: these four are not the consensus. Plenty of serious researchers still expect scaling plus refinement to keep delivering, and the paper's own limitations concede that surveying four critics risks making the field sound more settled against scaling than it is.

Useful and overhyped are not mutually exclusive

There is a strong pull, in any argument about this, to pick a side: either the technology is real and the money is fine, or the money is mad and the technology is a trick. Both are more comfortable than the position the paper actually takes.



Useful and overhyped are not mutually exclusive.

The sentence the paper flags as the most important one in its section on historical parallels

A technology can be genuinely important and the financial structure built around it can still have run ahead of what it has been shown to return. Those are two claims about two different things, and holding both is not fence-sitting.

Three cracks

None of these says the bet fails. Each is a place where the structure is thinner than the spending assumes.

THE ADOPTION GAP

The distance between organisations trying AI and organisations running something important on it. Pilots are plentiful. Production is not.

WHO IS MAKING MONEY

Most revenue so far belongs to the companies selling the equipment rather than the companies using it. One leading AI firm is reported to have lost around five billion dollars in 2024.

DEMO TO DEPLOYMENT

A system that is ninety per cent right is a remarkable demonstration. For work where being wrong carries a cost, the remaining ten per cent is the entire problem.

The loss figure comes from business journalism rather than audited accounts: an informed estimate, not a verified fact.

How much of the market this now is

The reason a wrong bet would not stay with the people who made it is arithmetic rather than conspiracy.

about 30%

of the total value of the S&P 500 sat in seven large technology companies, as of late 2025.

S&P Global, 2025. A point-in-time figure that moves with the market; quoted for order of magnitude, not as today's number.

roughly $\frac{3}{4}$

of the index's returns since ChatGPT are attributable to a small set of AI-exposed companies.

S&P Global, 2025. The paper calls this directional rather than measured: published figures run from about 60 per cent to over 80, depending on window, index method, and what counts as AI-related.

about 60%


of American workers are in an employer-sponsored retirement plan, most holding broad market funds.

Investment Company Institute, 2024.

Put those together and a decline concentrated in those companies does not stay in those companies.

Four ways the cost travels

Corrections are painful. The claim is that the pain lands somewhere specific, and not where the decisions were made.



The one in your account

A target-date fund is the most sensible, least adventurous thing a person can do with retirement money. It is also, right now, a concentrated position in a handful of companies, arrived at by following advice that was good advice.

The people who made this bet deliberately have options an ordinary saver does not: they can spread the risk, insure against it, and get out at a time of their choosing.

So the group that made the decision and the group carrying the consequence are largely different groups of people. That is the paper's central distributional claim, and it does not require anybody to have behaved badly.

The towns that built for it

Data centres land in particular places: northern Virginia, central Oregon, west Texas, parts of the Midwest. The construction work is real and it is temporary. The permanent staffing of a finished facility is small relative to how much ground it covers. Meanwhile local governments issue bonds and plan budgets against the tax revenue those facilities are expected to produce.

Around 2000 the same thing happened with fibre-optic cable: towns built for traffic that did not come, and were left with unused capacity in the ground and budgets planned around growth that never arrived.

The current build is larger, and it is more concentrated in fewer places.

Nobody announces this one

The way work actually disappears is not a press release about artificial intelligence. It is a role that does not get refilled when somebody leaves, a team restructured into a smaller one, a set of tasks that quietly moves.

WHAT IS UNSETTLED

Whether AI ends up destroying more work than it creates is genuinely open among economists. The paper cites both sides and resolves neither. The data is too young.

WHAT IS NOT

That the transition costs something, that it costs it to particular people in a particular decade of their lives, and that aggregate figures hide exactly that.

The author writes himself into this section rather than observing it from outside: after twenty-five years in technology consulting, made redundant in September 2023, then roughly a year looking for the same kind of job before deciding to do something else. He notes he had savings, a family, and a background that made the pivot possible, and that most people in that position have none of those.

The slowest channel

The fourth route is the one with no obvious victim, which is why it tends to go unmentioned.

COMPLEXITY

When the only people who understand a system work for the companies building it, those meant to govern it become dependent on them.

SPEED

Capability moves in months. Legislation moves in years. The result is not deliberate deregulation; it is deregulation by default.

WHICH LEVEL DECIDES

Where national rules override state and local ones, the places carrying the local consequences have the least say over them.

ATTENTION

The public has a limited amount of it and this subject is tiring. That is a problem with how the information reaches people, not with the people.

None of this is corruption, which is exactly what makes it hard. Each one makes the next harder to fix, and between them they leave a gap filled by the organisations that most needed governing.

If not this, then what

The paper does not stop at the diagnosis, but it is modest about the alternative, and so is this slide.

Shift the aim from building something that replaces the person to building something that makes a person's judgement go further. That is also a cheaper thing to build, because it does not require solving general intelligence first.

A companion piece, What Would You Build Instead?, sets out what such a system looks like and what it would commit to. The paper itself is careful to call this a direction rather than a specification, and says plainly that its author would need collaborators to turn it into one.

What this does not claim

Compression is where a careful argument turns into a loud one. So, explicitly.

✗ ~~A crash is coming.~~

Not the claim, and not a prediction. The stated purpose is to identify where the structure is fragile and who is hurt if it gives.

✗ ~~AI is a fraud, or useless.~~

The technology is treated throughout as real and important. Fragility is a claim about the financial structure on top of it.

✗ ~~The experts agree scaling has stalled.~~

They do not. Four critics are surveyed, and the paper concedes four critics are not a consensus.

✗ ~~AI is destroying jobs.~~

The economics is unsettled and the paper leaves it unsettled. The narrower claim is that the transition has costs that fall on identifiable people.

✗ ~~The author has no stake in this.~~

He states the opposite: displaced from the industry himself, and flags that this may incline him to read the signals as more threatening than they are.

Where these numbers come from

No other piece in this series has this slide, and this one needs it. The figures are not all the same quality, and the paper says which are which rather than letting the strongest-sounding ones carry the argument.

ADOPTION

Official statistics put business use of AI under about 10 per cent on a narrow definition. Private surveys using broader ones reach upward of 50. The paper uses the low figure deliberately.

CONCENTRATION

The three-quarters-of-returns figure is directional. Published estimates run from about 60 per cent to over 80 depending on method.

COMPANY LOSSES

Revenue and loss figures for private AI companies come from business journalism, not audited accounts. Informed estimates, not verified facts.

Three more in the paper: the 1.5 trillion dollar bond projection is a forecast rather than an observation, the labour question is genuinely open in the literature, and the energy figure comes from equity research rather than disclosures.

Three questions worth asking out loud

1. WHO DECIDED, AND WHO IS HOLDING IT?

If those are the same people, this is just investing. If they are different people, it is something else, and the difference is the whole subject.

2. WHAT WOULD HAVE TO BE TRUE FOR THIS TO WORK?

Ask it of the four steps. A chain where everything rests on one claim is not four risks. It is one risk, four times.

3. WHERE IS THE NUMBER FROM?

Survey or census. Audited or reported. Measured or projected. The difference between a figure that can bear weight and one that cannot.

Where this comes from

This is a short version of a working paper called When the Music Stops, which carries the full argument, the sources, the limitations and the open questions, including the ones the author cannot answer.

Part of a series on what judgement is, and why it has to be carried by people who bear the consequences of being wrong.

If we do not choose the architecture, the architecture will choose us.

Working papers are revised when they are shown to be wrong, and what changed is kept on the record.

James Maconochie | Architecture & Attention

HUMILITAS, HUMANITAS, ET CURIOSITAS · SEMPER

Humility, humanity, and curiosity. Always.

