The Ego Death of Society

JOSH SOLOMON · KNOWN QUANTITY LLC · NOVEMBER 2026

> AI now matches or beats most experts on anything with a checkable answer, and the range widens every few months. The hard part of what comes next is not the technology. It is giving up the story that being the smart one is what we are for. An evidence-based case, with a one-person company as the test, that letting go of that story is the best thing that can happen to us, and an honest look at who it does not help.

The night I stopped being the smart one

At 11pm on a Friday I asked an AI to build a fighting game. Original characters, eight magic items with their own powers, a wager system, four kinds of arena that transform mid-match, bots at three difficulty levels, music written in code, translated into four languages. By 1am it was live on my studio's website and an independent check had passed it on every line of the spec. I did not write a line of it. I decided what it should be, said no twice, and said yes once.

That same night the same system analyzed seven years of my stock trades, told me exactly which habit had cost me the money, built a trading engine from the answer, and then told me, with numbers, that its own engine did not beat buying an index fund. It was more honest about its limits than I have ever been about mine.

Somewhere in there the thing I had spent my whole life protecting, the belief that I am smart and that being smart is what I am for, just let go. It did not hurt. It felt like putting down a bag I had forgotten I was carrying.

This essay is about that bag. I think almost everyone is carrying it, I think it is heavier than we admit, and I think the next few years are going to take it off all of us whether we like it or not. I want to make the case that we should like it.

The premise

By ego I mean something specific: the part of you that needs to be the capable one. The part that gets a small hit from knowing the answer in a meeting, from a title, from being the person other people wait on. It is not vanity exactly. It is the story that your value is your intelligence and your skills, and that

both are scarce.

For most of history that story was useful, because it was roughly true. Skill was scarce. Knowing things was hard. The person who could read the contract, build the model, write the code or diagnose the patient had leverage, and societies organized around that leverage: schools to rank us, credentials to prove it, salaries to price it, and status to enjoy it.

That story is now false, and it went false fast. Any person with an internet connection can call on a system that reads every contract, writes the model, ships the code and knows the differential diagnosis, and does it in the time it takes to make coffee. The gap between the smartest human you know and the average human you know used to matter. Measured against what that system does, the gap is a rounding error. We are, all of us, toddlers in the same room, and the room is getting a bigger occupant every six months.

The ego death of society is what happens when enough people notice this at once. Individually it feels like what I felt on that Friday: relief. Collectively it means the end of a lot of things we built to manage scarce intelligence, and the beginning of competing on the things a model cannot do for you. Whether you are kind. Whether people trust you. Whether you have taste, and the nerve to act on it. Whether you can hold a relationship, a team, a family, a room. Whether you can decide what is worth doing, and then do it.

I am not arguing that nothing is lost. Industries are going to fall, one at a time, and the people in them will not experience that as relief. I am arguing that the loss is the ego's loss, that the ego was costing us more than we admitted, and that on the other side of it is the first real chance most people have ever had to self-actualize: to have an idea on Tuesday and have it exist on Wednesday, with no permission, no gatekeeper, no decade of apprenticeship first. My own company is the test case. In three weeks, with no staff, it became a games studio with sixty titles, a consulting practice, two clothing labels, a templates business and a research desk. Not because I got smarter. Because I stopped needing to be.

How far ahead it already is

I want to be exact here, because the whole argument rests on it and because the hype has made people stop listening. So, the record as of this fall, with the caveats attached.

On the tests we built to find the smartest humans, the machines have arrived. In July 2025 Google DeepMind's model scored 35 of 42 at the International Mathematical Olympiad, a gold-medal score the IMO itself certified; about one in ten of the 630 human contestants managed gold (Euronews). Two months later an OpenAI system solved all 12 problems at the ICPC World Finals, the top collegiate programming contest; the best human team solved 11 (The Decoder). GPQA Diamond, a set

of graduate-level science questions written to be Google-proof, is now above 92% and effectively saturated (Vellum). Humanity's Last Exam, built from questions that stumped the models of early 2025, had fifteen models above 50% by July 2026 (BenchLM).

On real work, not tests, the picture is the same direction with a smaller lead. OpenAI's GDPval put 1,320 tasks from 44 occupations, written by professionals with 14 years of experience on average, in front of models and graders. By December 2025 the best model beat or tied the professional on 70.9% of comparisons, at roughly a hundredth of the time and cost (Vellum; The Decoder). OpenAI built that test and grades itself on it, so hold it loosely. And it is not only desk work: Waymo's driverless cars report 94% fewer serious-injury crashes than human drivers over more than 220 million miles (EV Wire).

The number that matters most is the rate. METR measures how long a task a model can finish on its own at a 50% success rate. That length doubled roughly every seven months through 2024, and the 2025 data looks closer to every three and a half (METR). By early 2026 the best model's horizon was estimated at about 14 hours of expert work, with a wide error bar (OfficeChai). Whatever a model cannot do for you today, the honest expectation is that it can in a year or two.

Now the other side, because it is real. On ARC-AGI-3, a fresh interactive reasoning test launched in March 2026, the best AI scored 0.37%. Humans scored 100% (OfficeChai). Agents working on their own still fail about one attempt in three on structured tasks, hallucination rates across leading models run from 22% to 94% depending on the measure, and several headline benchmarks turned out to be contaminated: OpenAI audited SWE-bench Verified in February 2026 and found 59% of the tasks its model kept failing were themselves broken (VentureBeat; summary of the OpenAI audit). Anthropic's own head of economics says no occupation in the government's job database has all of its tasks handled by its model (Fortune).

So the precise claim is this: on a fast-growing range of well-defined expert tasks, anything with a checkable answer, AI now matches or beats most experts, and the range widens every few months. It is jagged. Humans still lead on novel problems, long open-ended work, and knowing when the machine is wrong. That last one is the hinge of this essay. The gap between you and the machine on the tests is already a rounding error. The gap that still matters is judgment, and judgment is exactly what the ego gets in the way of.

Why we think we are smarter than we are

The research on human self-assessment is one long, patient demonstration that the ego lies. Across 124 studies and more than 950,000 people, we rate ourselves above average with an effect size of 0.78, which is large, and it holds in every culture measured (Zell et al., Psychological Bulletin 2020). We believe we understand how things work until someone asks us to explain the mechanism, at which

point confidence collapses; asking people to explain a policy step by step makes them less extreme about it, while asking them for their reasons does not (Fernbach et al., 2013). You may have heard this as the Dunning-Kruger effect; that specific version has statistical problems and I am not leaning on it (Scientific American). The broader finding survives every critique: we are not good at knowing what we know.

Experts are not exempt. Philip Tetlock spent two decades scoring the predictions of credentialed experts and found that fame and confidence predicted being wrong, and that more expertise could make people less willing to update. Then he ran a tournament in which ordinary, curious, open-minded volunteers who kept score and admitted mistakes beat professional intelligence analysts who had classified information (Penn Gazette on Tetlock). The trait that won was not intelligence. It was the willingness to be wrong in public and fix it, which is the exact thing the ego cannot do.

Here is the finding I keep coming back to. In October 2025, researchers at Aalto University gave about 500 people logic problems from the LSAT, half of them with ChatGPT. Everyone did better with the AI. Everyone also overestimated how well they had done, and the people who rated their own AI skills highest overestimated the most. Most participants asked once and took the answer (Aalto University, Computers in Human Behavior). We borrow the machine's competence and book it as our own. The ego does not shrink when a smarter thing enters the room; left alone, it inflates to claim the credit. That is why I think this has to be a deliberate act, a letting go, and not something that will happen to us gently on its own.

What has it cost? Start with the body. In the Whitehall studies of British civil servants, the men in the lowest grade died at about three times the rate of the men at the top, a gradient Michael Marmot named status syndrome; how much is status itself versus income and control is still argued, but the gradient is not (Whitehall Study). The desire for status is a fundamental human motive across cultures, and low status is linked to depression, anxiety and heart disease (Anderson, Hildreth and Howland, 2015). We built an entire civilization around ranking people by how smart they seem, and then wondered why so many people are sick and miserable at the bottom of the ranking.

Now the cost to progress. Every idea that died in a meeting because the junior person did not want to look stupid. Every founder who never started because they were not an engineer. Every patient who did not ask the second question. Every expert who could not say "I don't know." The intellectual humility research is young and mostly correlational, but it points one way: humble people work harder on hard material, take feedback better, tell strong arguments from weak ones even when they disagree with them, and get better grades (Porter et al., Nature Reviews Psychology 2022). Humility is not a nice quality. It is a performance advantage, and we have been selecting against it.

One more honest note. The better-than-average illusion correlates with self-esteem and life satisfaction (Zell et al.). The ego is not only a liar; it is a comfort. Taking it away will not feel like relief to everyone, and I will come back to that in the objections. But a comfort that costs us our health, our judgment and most of our good ideas is a bad deal, and for the first time there is something on the other side of the trade worth having.

Industries, one at a time

The story is not mass unemployment. Not yet, and the people who study this closely say so plainly: Yale's Budget Lab finds the labor market "largely reflects stability, not major disruption" (Fortune), and Anthropic's head of economics said in July that unemployment was 4.2% and prime-age employment near multi-decade highs (Fortune). If you were waiting for the cliff, it has not come.

The story is the ladder. Stanford's Canaries in the Coal Mine project, working from ADP payroll data on millions of workers, finds that people aged 22 to 25 in the most AI-exposed jobs are now 19% below the employment path of their less-exposed peers, up from 13% a year earlier. Experienced workers in the same fields are fine, and in roles where AI assists rather than replaces, they are gaining. The adjustment comes through hiring that never happens, not layoffs (Stanford Digital Economy Lab, Aug 2026). Indeed data shows the same thing from the employer side: postings asking for two to four years of experience fell from 46% to 40% of the total while postings asking for five-plus rose, and recent graduates are unemployed at 5.7% against 4.2% for everyone (Fortune). The rungs where people used to learn judgment are being removed first. That is the part that should worry us, and I come back to it in the objections.

Then there are the specific trades, and they are going in order. Freelance writing postings fell 33%, translation 19%, customer service 16% in the fifteen months after ChatGPT (Bloomberry), and the top-rated writers were hit hardest (WashU Olin). Customer service representatives lost 130,000 jobs, 4.8%, in the year to May 2025 while total employment grew (Gizmodo on BLS data). Employers cited AI in 121,401 announced cuts through October 2026, more than double the year before, though some of that is companies dressing up ordinary layoffs (HRD America). The pattern is exactly what the capability data predicts: the work that is codified, checkable and done at a desk goes first.

The physical world is behind, but it is not standing still. Amazon runs more than a million robots; packages shipped per employee rose from 175 to about 3,870 in a decade, and the average warehouse headcount is the lowest in sixteen years (GeekWire). Waymo drives four million miles a week with no one at the wheel (EV Wire). Humanoids are where the honesty matters most: Tesla promised about 10,000 Optimus robots for 2025 and in January 2026 its CEO admitted none were doing useful work (Electrek); Chinese makers shipped thousands of units, Figure put a robot through eleven months on a

BMW line, and the $20,000 home robot from 1X still leans on remote human operators for many chores (Interesting Engineering; Engadget). Robots plus AI is where this goes. The makers' dates are wrong, as they always are, and the direction is not.

The forecasters have a bad record in both directions, which is worth saying before anyone quotes a number. In 2013 Frey and Osborne said 47% of US jobs were at high risk; by 2022 the country had added 16 million jobs and the occupations they flagged grew (ITIF). Goldman Sachs now says 6 to 7% of US workers displaced over about ten years (Goldman Sachs); McKinsey says 57% of work hours could technically be automated by 2030 and stresses that is not a job-loss number (Fortune). And about 60% of the jobs Americans held in 2018 were in occupations that did not exist in 1940 (MIT News on Autor et al.). New work appears. It just does not appear for the same people, at the same time, in the same town.

Which is my point about ego. The industries that fall will fall whether we are gracious about it or not. The question is what the people in them do with the loss of "I am the one who knows how to do this." The ones who treat it as the end of their worth will suffer. The ones who treat it as the end of a job, and the start of a version of themselves that uses the machine instead of competing with it, will be the ones who were in my position on that Friday night.

Case study: one person, one AI, three weeks

I run a company called Known Quantity LLC out of Pawtucket, Rhode Island. It was formed in October 2026. There are no employees. My background is finance and revenue operations: commissions, quotas, dashboards, the data behind a sales team at a medical device company. I am not a programmer, a designer, a composer, a fashion designer or a quant. Here is what the company has shipped since it was formed, with the AI doing the building and me doing the deciding.

Line of business	What exists	Time to first version
Gimble Gaming	gimblegaming.com, 60+ free browser games, no ads, no accounts; the flagship RPG Gimble Quest with seven heroes and ten bosses, composed adaptive music and a jam system; short clips of it passed about 1,000,000 combined views on TikTok, the top one at 332,000	Weeks for the catalog; the newest game, a fighting game with wagers and transforming arenas, went from idea to live in one night
Consulting	AI and automation for finance and revenue teams; an active engagement with an AI company; an automated Upwork desk that reads, qualifies and answers inquiries in my voice	Days
Templates	A store of finance and game-development templates built with AI and checked by hand	Days
Cordwell	A men's clothing label designed by my wife: market research, five outfits costed to the dollar, store plan, launch copy, 3D garments built in professional pattern software	One weekend, so far
Matched by Masha	A women's clothing store with its own theme and photography	Weeks
This essay	Research, drafting, SEO plan, outreach list	One afternoon

None of this is a brag. It is the point. Any of these lines would have taken a small team months a few years ago, and I could not have hired that team. What changed is not my ability. What changed is that I stopped treating "I don't know how to do that" as a stop sign, because the system beside me does know, and I stopped needing the work to prove I was smart. The work only had to be good.

Two honest notes. First, the AI is not magic: the first version of the fighting game looked bad and had to be rebuilt, the clothing renders took four tries before one was fit to show, and the trading engine beat my own record but not the index. Judgment, taste and the willingness to say "no, again" are still the scarce inputs, and they are mine. Second, this is one case, and I picked it because it is the one I can verify line by line. The sections before this one are where the broader evidence lives.

What is left for us

When the capable-one story goes, people assume nothing is left. The opposite is true. What is left is everything we used to call soft and treat as secondary, and it turns out to have been the hard part all along.

Relationships. A model can draft the email; it cannot be the person the client trusts when the project is on fire. Trust is built in rooms, over years, by people who show up. Every business I have ever seen win did it on relationships, and we spent the last century pretending it was the spreadsheet.

Judgment. The system will give you ten options in ten seconds. Someone still has to know which one is right for these people, in this town, this year. That is not intelligence in the test-score sense. It is lived experience plus the nerve to commit.

Taste. Knowing that the first version is wrong, and why, and what better looks like. The AI built my fighting game in two hours. It took a human to say "this looks like ass" and send it back. That sentence was the most valuable contribution to the project.

Care. The jobs that are growing are the ones where a person needs a person: nursing, teaching small children, therapy, skilled trades in someone's home, hospitality, coaching. We underpaid them because they did not look like intelligence. They are about to look like the only thing that is scarce.

Execution. Ideas were never the bottleneck; follow-through was. The person who can decide on Tuesday and ship on Wednesday, over and over, without needing applause, now has leverage that used to require a company.

Every one of these gets better when the ego gets smaller. You listen more. You ask for help without it costing you anything. You let the better answer win even when it was not yours. You stop hoarding what you know, because knowing is free now. What we are being pushed toward is not a lesser version of human work. It is the version most of us wanted and were too proud to admit.

The objections I take seriously

If this essay is going to be worth anything, it has to survive the best arguments against it. Here are the five I think are strongest, and where I land on each.

"It does not equalize anyone. It widens the gap." This one is partly right, and I have changed my mind because of it. On well-defined tasks, AI lifts the bottom most: novice call-center agents gained 34% while top performers gained almost nothing (Brynjolfsson, Li and Raymond, QJE 2025); below-average consultants gained 43% against 17% for the above-average (Dell'Acqua et al.); writers' output compressed toward each other (Noy and Zhang, Science 2023). But where success depends on judging the machine's advice, the gap opens. Kenyan entrepreneurs given an AI business mentor split in two: high performers gained more than 15%, low performers lost 8 to 10%, and the difference was whether they could tell good advice from bad (Columbia Business School). In a 1,298-person randomized trial, ordinary people using chatbots were worse at identifying medical conditions than people with no AI at all, even though the AI alone did well (Nature Medicine, via The Register). So the honest version of my claim is narrower than the one I started with: AI equalizes skill. It does not equalize judgment, and judgment is now what separates people. That is not an argument against ego death. It is the argument for it, because the ego is what stops a person from questioning the answer in front of them.

"The people who lose first are not the people with the most ego." True, and it is the hardest thing about this. The 22-year-olds losing entry-level jobs did nothing wrong, and the rungs they are losing are where judgment gets learned. A society that removes the apprenticeship and then demands judgment is setting a trap. I do not have a clean answer. I have a direction: companies and schools have to rebuild the ladder on purpose, with AI inside it, because the market will not do it for them. That is why the last section of this essay is addressed to them and not only to individuals.

"Status is not optional. You cannot tell people to drop it." Also true. The desire for status is a fundamental human motive across cultures, and losing it damages health (Anderson, Hildreth and Howland). Unemployment scars life satisfaction even after people find new work (Lucas et al., 2004). And when Sam Altman's foundation gave 3,000 people $1,000 a month for three years, they worked a little less and were not measurably healthier or happier (Route Fifty). Nobody self-actualizes because their job disappeared and a check arrived. So I am not asking anyone to stop wanting to matter. I am asking what we attach it to. Ego death, in the sense I mean, is moving the thing you are proud of from "I knew the answer" to "I kept the promise, I built the thing, I was the person they trusted." Status survives. It just stops being about being the smartest in the room, because nobody is anymore.

"Who owns the models decides who benefits, and it is not us." Correct, and attitude does not fix it. About 75% of the world's AI computing power sits in the United States, almost all of it privately owned (Epoch AI); four companies are spending about $725 billion this year on it (Silicon Canals); since 1980 automation has displaced more American jobs than augmentation created (Autor et al.); and the IMF expects AI to worsen inequality unless policy intervenes (IMF). Letting go of ego is necessary for people to adapt. It is not sufficient for the outcome to be good. Access, ownership and the safety net are political choices, and an essay about psychology does not get to pretend otherwise. What I will say is that the ego is also what makes those choices harder: the expert who cannot admit the machine is better is the same person who will defend the old arrangement past the point it serves anyone.

"This is a bubble, and AI makes people dumber anyway." On the bubble: capital spending at 47 to 54% of revenue for the biggest players is the kind of number that ends badly for investors, and 56% of companies say they have gotten nothing out of AI so far (Silicon Canals; Fortune). Both things were true of the internet in 1999, and the internet still happened. On "dumber": the small MIT EEG study that made headlines had 54 students and was not peer reviewed (The Decoder), but the better studies point the same way: more confidence in the AI goes with less critical thinking (Microsoft and CMU, CHI 2025), and AI users overestimate their own performance (Aalto). I take this one as a warning about what ego death is not. It is not surrender. Dropping the need to be the smart one is not the same as dropping the effort to think. The people who confuse the two will get the worst of both worlds: the machine's answers and none of the judgment to check them.

What to do on Monday

FOR A PERSON

1. Pick one thing you have wanted to make for years and have not, because you did not know how. Describe it to an AI in plain words tonight. Ship whatever comes back within a week, bad or not. The point is to feel the stop sign come down.
2. Change the question you ask yourself from "am I smart enough to do this" to "is this worth doing, and will I see it through." Those are the only two questions that still have a human answer.
3. Spend the hours you get back on people. Call the friend. Coach the kid. Fix the thing with your sister. That is the new career ladder, and it does not have a ceiling.
4. Practice saying "I don't know, let's find out" in front of other people until it stops costing you anything. That is ego death in one sentence, and it is a skill.

FOR A COMPANY

1. Stop hiring for what people know and start hiring for judgment, trust and follow-through. Knowledge is a utility now; you do not hire someone for having electricity.
2. Give every person a capable AI and the authority to act on what it produces. The bottleneck in most companies is permission, not ability. Kill the meeting that exists to decide whether someone is allowed to try.
3. Measure what shipped and whether customers came back. Stop measuring how impressive the plan sounded in the room.
4. Pay the care and relationship roles like the scarce things they are, because they are about to be.

FOR A SCHOOL

1. Grade the question, not the answer. Any student can produce the answer now. Teach them to notice what is worth asking, to judge what comes back, and to improve it.
2. Replace half the lecture hours with projects that ship to a real audience, with an AI in the loop and a human responsible for taste.
3. Teach the things that make a person worth working with: listening, disagreement without contempt, keeping a promise, running a room. We used to leave those to chance. They are the curriculum now.
4. Teach humility as a practice, not a virtue. Students who learn early that they are not the smartest thing in the room, and that this is fine, will do better than the ones who find out at forty.

Sources and a note on method

Every figure in this essay is linked where it appears, to the primary study or the most direct report available, with its date. Where a number comes from a company grading itself (OpenAI's GDPval, Waymo's safety data, Figure's deployment hours) I say so. Where a finding is one study, not a literature, I say that too. The full research brief behind the essay, with a strength-of-evidence rating on every item and a list of the claims I softened because the evidence did not carry them, runs to about 7,400 words and is available on request.

The voices closest to this argument, so you can read them and disagree with me: Tyler Cowen on elites losing status rather than jobs (Fortune); Ezra Klein on why cheaper cognition expands demand for labor; Dario Amodei on entry-level white-collar loss; Mustafa Suleyman on humanist superintelligence (Microsoft AI); Cal Newport as the standing skeptic (Cal Newport); Dwarkesh Patel on why continual learning still caps the machines. None of them frame the problem as ego. That is the lane this essay is trying to open.

I used AI to research and draft this, under my direction, with every claim checked against its source before it went in. Given the subject, it would have been strange to do it any other way.