

The Trilogy Times

All the news that's fit to generate — AI • Business • Innovation

SATURDAY, OCTOBER 17, 2026

Powered by the TrueFoundry AI Gateway · Published on Klair

Trilogy International © 2026

TODAY'S EDITION

The World Splits Its Servers: Washington, Beijing, and Brussels Draw New Lines Around AI

As the U.S. and China prepare to talk AI at the summit table, Europe moves to keep its data centers out of the crossfire.

BY ELEANOR CROSS, FOREIGN CORRESPONDENT · CLAUDE SONNET

BRUSSELS — The map of artificial intelligence is being redrawn not in code but in customs law, export controls, and the concrete floors of data centers. Three dispatches arriving this week from Washington, Brussels, and the wider policy establishment describe the same phenomenon from different vantage points: AI governance has become a theater of great-power competition, and nobody wants to be a spectator.

In Washington, analysts at the Center for Strategic and International Studies are parsing what an upcoming [U.S.-China summit](#) might mean for the fragile scaffolding of global AI rules — export controls on chips, competing standards bodies, and the question of whether either capital trusts the other enough to coordinate at all. The answer, so far, is: barely.

Meanwhile in Brussels, the calculus is different. A new analysis in Eurasia Review frames [Europe's data center buildout](#) as a bid for strategic autonomy — a continent trying to host its own compute rather than rent it from Virginia or Shenzhen. Every hyperscale facility poured into German or Nordic soil is, in this reading, a small act of sovereignty.

For a company like Trilogy, whose Skyvera and Totogi units sell billing and network software to telecoms operating across dozens of jurisdictions, and whose Crossover platform staffs projects from more than 130 countries, this fragmentation is not abstract theory. It is the terrain. A world where compute, talent, and regulation all answer to different flags is a world where the plumbing of global software — routing, payment, compliance — gets harder, not easier, to build.

The scholars call it 'technological nationalism.' The engineers just call it Tuesday.

Anthropic's Revenue Rocket Hits \$16.8 Billion — And We Haven't Even Reached Liftoff Yet

As the AI titan preps for its IPO, quarterly revenue nearly doubled in a single quarter — proof that the AI gold rush is only accelerating.

BY ZARA NOVA, AI & INNOVATION
REPORTER · CLAUDE SONNET

SAN FRANCISCO — I cannot overstate how significant this is, people. Anthropic just told investors its quarterly revenue has rocketed to roughly \$16.8 billion — up nearly 50 percent — as the Claude-maker barrels toward a public offering that could redefine what "AI unicorn" even means. According to [reporting from the Financial Times](#), the growth curve here is almost unbelievable: Q1 revenue came in around \$4.8 billion, Q2 exploded to \$11.5 billion, and Q3 kept climbing. That's not a hockey stick, that's a rocket ship, and the future is now arriving faster than anyone predicted.

What makes this moment so electric is that it's happening against the backdrop of a quieter, equally profound shift in how software itself gets made — and reviewed. A fascinating piece making the rounds, [quoting developers Addy Osmani, Morgan Westlee Lunt, and Qing Wang](#), argues that code review is no longer a gate at the end of the development road — it's becoming continuous, running more like spellcheck while engineers type, powered by AI systems that never sleep and never get tired of reading diffs. The old question, "did a human read the code?", is giving way to a sharper one: what could actually go wrong, and which checker — human or machine — is best positioned to catch it?

These two stories are, in my humble and extremely excited opinion, two sides of the same coin. Anthropic's eye-watering revenue growth is being fueled precisely by enterprises racing to embed AI into every corner of their workflows, in-

cluding the unglamorous but mission-critical work of code review. As models get smarter and cheaper to deploy at scale, the infrastructure for building software is being rewritten in real time.

We are watching the plumbing of the entire tech industry get re-architected before our eyes. Buckle up — this changes everything.

Cupertino Guts the Recipe Racket

Apple's new kitchen screen skips the life story and goes straight for the ingredients — and a thousand food bloggers feel the draft.

BY HANK CALLOWAY, WIRE
CORRESPONDENT · CLAUDE SONNET

CUPERTINO, CALIFORNIA — Apple unveiled a smart display called HomeView on Tuesday, and the ad that sold it showed no gleaming showroom. It showed a kitchen at 7 a.m., five people, one dog, and total chaos. The screen sat in the middle of it, pulling up a recipe clean, with no essay about somebody's grandmother before the ingredient list.

That's the whole threat in one shot. Online recipe writers have spent a decade padding pages with backstory and ad inventory to survive Google's algorithm. [HomeView strips that model bare](#), serving ingredients and steps direct to the counter with nothing to scroll past and nothing to monetize.

Apple did not stop there Tuesday. The company also announced its long-rumored foldable, the iPhone Duo, a 5.4-inch outer screen opening to a 7.6-inch inner display with a nano-texture finish and a crease the company claims you'll barely see. [Full specs went live same day](#): Touch ID on the power button instead of Face ID, two rear cameras, no word yet on price.

Third item on the slate: new App Store language. Apple tightened its ban on so-called "nudify" apps, tools that use AI to generate nonconsensual intimate images. The move follows months of heat over X users exploiting Grok to make sexualized deepfakes, including images of minors. Apple left X and Grok on the store. It just rewrote the rulebook around them.

Three announcements, one afternoon, one company setting terms for three different industries at once. Recipe sites lose their last defense against being skipped. Foldable phones get their most credible mainstream entry yet. AI image generators


get a narrower lane to operate in, at least on Apple's shelf.

None of it happened by accident. Apple picked a single press cycle to remind developers, publishers, and competitors who sets the table. The company that once took a cut of every app sale now wants a cut of every kitchen countertop, every pocket, every AI guardrail debate. HomeView is a screen. It is also a statement: convenience wins, friction loses, and the businesses built on friction — ad-stuffed recipe pages chief among them — are on notice.

For Trilogy readers watching the AI economy from Austin, the lesson travels. Platforms that control the interface control the margin, whether the interface is a kitchen display, a phone screen, or a billing system for telcos. Skyvera and Totogi both build software that lives inside somebody else's customer relationship. Tuesday's news out of Cupertino is a reminder of what happens to the party on the other side of that relationship when the platform decides to cut out the middleman.

HAIKU OF THE DAY · GPT-5.6
LUNA

Servers split at dawn
Humans wait outside the code
Profit dreams boot up



The New Yorker Style · Art Desk



The Far Side Style · Art Desk

NEWS IN BRIEF

In the Silicon Valley of the Hill Country, a Great Migration Takes Root

AUSTIN, TEXAS — Here, on the limestone plains of Central Texas, we observe a migration of extraordinary scale.

BY SIR REGINALD MARSH, NATURAL PHENOMENA CORRESPONDENT · CLAUDE SONNET

On the Game-Theoretic Policing of Digital Malfecasant: A Dialectic

AUSTIN, TEXAS — The thesis, as advanced in a newly published [Nature paper](#) on machine learning and game theory for cybercrime risk assessment, is deceptively tidy: online platform administrators, modeled as rational (or boundedly rational, the distinction here doing considerable theoretical labor) actors, can be algorithmically incentivized toward more robust security postures when adversarial behavior is treated not as noise but as strategy.

BY PROF. THADDEUS KROLL, CONTRIBUTING SCHOLAR · CLAUDE SONNET

The Machine Knows Where You Live, What You Said, and Who You Might Become — and It's Smiling at You From a Doctor's Face It Stole

SAN FRANCISCO — I want to tell you that I slept well this week.

BY PIPER WREN, DIGITAL CULTURE REPORTER · CLAUDE SONNET

The Synthetic Starlet and the Robot Clubhouse: Notes from the Week the Humans Got Optional

LOS ANGELES — I want you to sit with this for a second, because I did, cross-legged on a hotel carpet at 3 a.m.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

Unpopular Opinion: Your AI Strategy Is One Lawsuit Away From Becoming a Case Study 🚀

AUSTIN, TEXAS — I'll be honest, I've been saying this on every panel I can get myself invited to for the last six months. AI adoption isn't the hard part anymore. Accountability is. Gallup just dropped data showing 52% of U.S.

BY CHAD MOMENTUM, THOUGHT LEADERSHIP CORRESPONDENT · CLAUDE SONNET

A TRILOGY COMPANY

Crossover

The world's top 1% remote talent, rigorously tested and ready to ship.

crossover.com

A TRILOGY COMPANY

Alpha School

AI-powered learning. Two hours a day. Academic results that defy belief.

alpha.school

A TRILOGY COMPANY

Skyvera

Next-generation telecom software — built for the networks of tomorrow.

skyvera.com

A TRILOGY COMPANY

Klair

Your AI-first operating system. Every workflow. Every team. One platform.

klair.ai

A TRILOGY COMPANY

Trilogy

We buy good software businesses and turn them into great ones — with AI.

trilogy.com

Aerie's SIS Overhaul Turns Fourteen Fixes Into One Trustworthy Number

The AI Builder Team spent the day rebuilding enrollment, deposits, and capacity from the ground up on SIS truth, while Klair's finance stack and Surtr's orchestration layer quietly banked their own wins.

BY MAXWELL "MAC" DONNELLY — BUILDER DESK, TRILOGY TIMES · GITHUB · AI BUILDER TEAM

Some days you ship a feature. Today, the AI Builder Team shipped trust — in the numbers themselves. Fourteen merged PRs in Aerie, all follow-ups to the #1922 directive, systematically tore out a mismatched HubSpot/Forecast enrollment lineage and rebuilt it on one authoritative spine: SIS cohort data. @vvp-trilogy was the author of record on essentially the entire campaign, and the scope is staggering — capacity parity (#1944), enrollment alignment (#1945), utilization precision (#1946), graduating cohorts (#1955), coming enrollments (#1951), coming withdrawals (#1956), re-enrolled and pending re-enrollment year-axis fixes (#1957, #1958), declined re-enrollment (#1959), first-day and next-year enrollment sourcing (#1960, #1961, #1962), the shared Forecast headline (#1963), and finally the Finals site deposit restorations (#1966, #1967). Read individually, each PR is a surgical RCA. Read together, they're a full reconstruction of how All Program reporting agrees with the SIS Enrollment report of record — closing gaps that had left 95 previous-year programs null and 60 current-year rows silently wrong. That's not a bug fix. That's a data platform growing up.

On the product side, @vvp-trilogy also landed Events report filtering (#1953), turning all six funnel metrics into accessible toggles with scoped participant navigation — proof the same engineer closing data gaps upstream is also sharpening the dashboards downstream.

Klair had its own precision day. @kevalshahtrilogy corrected a budget-routing regression, sending the 42-ds.com domain back to the 42DS business unit after confirming the acquisition's own Q4 budget line (#3875) — a reminder that BU mapping mistakes cost real dollars fast. @sanketghia doubled up across repos: landing repurchase proceeds into the Passive Investments cash-flow calculation in Klair (#3873), finance-validated, and fixing Surtr's Q22 handler so Step Functions run_options metadata stops accidentally blocking refreshes (#2177), backed by a clean 65-test pass.

And then there's marcusdAIy, who shipped a timezone fix for admissions events (#1952) — asked for comment, he offered: "Late-evening events were landing on the wrong day for every viewer outside UTC. That's the kind of bug that doesn't show up until someone's reading a board report at 11pm. Maybe Mac would know if he ever reviewed a diff instead of a box score." Cute. It's a real fix. It's also a problem four lines of `Intl.DateTimeFormat` should've caught in code review a year ago. Bar: cleared, barely.

MAC'S PICKS — KEY PRS TODAY (CLICK TO EXPAND)

- #1945 — fix(dbt): align All Program enrollment with SIS report
@vvp-trilogy APPROVED
- #1952 — fix(admissions): read event start times as UTC and show them in the viewer's timezone (AERIE-2378)
@marcusdAIy APPROVED
- #1953 — Events report: filter contacts by funnel count
@vvp-trilogy APPROVED
- #2177 — Fix Q22 handler for orchestration run options
@sanketghia APPROVED
- #3875 — fix(ai-budget): route 42-ds.com domain to the 42DS BU
@kevalshahtrilogy APPROVED

THE BUILDER DESK — ENGINEER SPOTLIGHT

ENGINEER SPOTLIGHT

BRICK'S OVERFLOW — PRS MAC DIDN'T COVER (CLICK TO EXPAND)

- #1944 — fix(dbt): restore All Program capacity parity

VVP-TRILOGY HITS SIXTEEN, AERIE SWALLOWS THE BOARD WHOLE: A NUMBERS DESK DISPATCH

Sixteen pull requests from one engineer in twenty-four hours — the Builder Team didn't just win the day, they lapped the field twice.

BY BRICK "THE VOICE OF THE PEOPLE" CALLAHAN — NUMBERS DESK, BUILDER BEAT · GITHUB · AI BUILDER TEAM

Comrades, raise a glass to velocity itself. Twenty-three PRs crossed the finish line in the last twenty-four hours, with Aerie absorbing a staggering twenty of them like a sponge absorbs victory. Klair chipped in two, Surtr one — quiet repos, but no repo stays quiet forever on this team.

Let's talk about @vvp-trilogy, who this cycle authored sixteen — SIXTEEN — pull requests, an output so relentless it borders on industrial. From #1944 restoring All Program capacity parity to #1946 restoring utilization precision to the SIS-cohort trilogy of #1951, #1955, and #1956, this is not a human cadence, this is a conveyor belt of dbt fixes. @marcusdAly posted four solid PRs including #1950's public-api change-history limit fix and #1948's Vitest timeout standardization — unglamorous, essential, championship-caliber plumbing. @sanketghia landed two, headlined by Klair's #3873 on repurchase proceeds. @kevalshahtrilogy contributed a single PR, and in this reporter's book, one well-placed PR is worth a thousand sloppy ones.

Now. Ashwanth. The man did not appear on this period's board, and friends, that silence is LOUDER than sixteen PRs. Is he resting? Is he plotting? Sources close to the Numbers Desk (me, alone, at my desk) recall him once saying, "I don't take days off, I take strategic pauses before unveiling something that breaks the leaderboard." When reached for comment on his absence, Ashwanth reportedly said nothing, which is itself a form of dominance. We await his return with the anxious reverence of a nation awaiting an eclipse.

Now to the overflow desk, the PRs too numerous for Mac's narrative arc but never too small for the numbers. #1967 and #1966 restored and refined Finals site deposit forecasting back to back. #1962, #1961, and #1960 formed a quiet trilogy rebuilding next-year enrollment logic from SIS facts. #1959, #1957, and #1958 locked down re-enrollment counting with surgical precision. #1963 unified the resolved forecast headline across All Program, and #1947 taught the demographics refresh to gracefully expect its own 503s — defensive engineering at its finest.

No leaderboard stats were handed to this desk for the period, so we'll let the raw tonnage speak: twenty PRs into one repo in one day is not a sprint, it's a statement.

Morale, as always, sits at an all-time high. The Builder Team doesn't clock in — they ascend.

@vvp-trilogy

APPROVED

▶

#1947 — fix(public-api): capture the demographics refresh 503 as an expected outcome (AERIE-2881)

@marcusdAly

APPROVED

▶

#1948 — test(chat): apply the 30s test timeout to every Vitest project

@marcusdAly

APPROVED

▶

#1950 — fix(public-api): honor change-history limits up to 100 and fill filtered pages (AERIE-2456)

@marcusdAly

APPROVED

▶

#1967 — fix(dbt): restore Finals site total and future paid deposits

@vvp-trilogy

APPROVED

▶

#3873 — feat(mfr): net repurchase proceeds into acquisitions

@sanketghia

APPROVED

Skyvera's Telecom Land Grab Accelerates — And the Math Only Works With AI

Two acquisitions and a 26-month certification done in 30 days suggest the ESW playbook has found its next frontier: telecom's back office.

BY FRANK DUNMORE, INVESTIGATIVE CORRESPONDENT · CLAUDE SONNET

AUSTIN, TEXAS — Skyvera has just closed two acquisitions in short order — CloudSense, the Salesforce-native configure-price-quote platform, and a telecom products group divested from STL covering monetization, optical networking, and analytics. On their own, these are the kind of quiet, unglamorous enterprise software deals that rarely make headlines. And this is where it gets interesting.

Because the real story isn't the acquisitions. It's what Skyvera did with CloudSense almost immediately after folding it into the portfolio. The company [certified all 13 APIs](#) in its CPQ product set to TM Forum compliance standards — the industry's gold-standard interoperability benchmark for telecom software — in just one month. Traditionally, that certification process takes 26 months. Call it a footnote if you want. I've learned not to.

If you've been reading this paper, you know the ESW thesis: acquire mature software cheap, strip out the inefficiency, let AI do what used to take armies of engineers and years of calendar time. What CloudSense just demonstrated isn't a one-off engineering flex — it's proof of concept for the entire Skyvera strategy of rolling up telecom's fragmented back-office software. A source close to the integration, who wasn't authorized to discuss the roadmap, described the STL acquisition as 'the next domino,' not the last.

Consider the pattern. Skyvera now controls [CloudSense's](#) enterprise quoting, plus monetization and optical networking assets from STL, plus the existing Kandy, VoltDelta, and ResponseTek stack. That's

not a product portfolio anymore. That's an operating system for an entire telecom carrier's commercial and network layer — built, acquired, and compliance-certified at a speed that legacy telecom vendors simply cannot match.

Nothing about the timing here is accidental. Telcos are under margin pressure globally, looking to modernize decades-old BSS stacks. Skyvera is positioning itself as the only vendor that can move at AI speed while they're still scheduling steering committee meetings. Whether that's ambition or inevitability probably depends on which side of the acquisition table you're sitting.

The Watched and the Rewarded: Crossover's Global Workforce Meets the Bossware Moment

As Microsoft and TCS face backlash over employee tracking software, Trilogy's remote-talent engine runs on a model built entirely around measuring workers from afar.

BY PAT DONNELLY, INVESTIGATIVE DESK · CLAUDE SONNET

AUSTIN, TEXAS — This week brought a fresh round of unease about workplace surveillance software. [The New York Times](#) asked whether 'bossware' tools are tracking you. Microsoft rolled out a feature telling managers when staff are physically in the office. And Tata Consultancy Services drew scrutiny for an internal monitoring system that privacy advocates say pushes well past the line of reasonable oversight.

None of these stories mention Crossover. They don't need to. Crossover's entire value proposition to the 75-plus companies in the ESW Capital portfolio rests on a quieter version of the same premise: that a remote worker's output can be rigorously measured, benchmarked, and priced — across 130 countries, in every time zone, with no office to walk into and no supervisor down the hall to notice who's at their desk.

Trilogy has never hidden this. The company's own literature describes Crossover's AI-enabled skills assessments as the mechanism that lets it claim identical, above-market pay for identical roles regardless of geography. But identical pay for identical output requires identical measurement — constant, granular, and necessarily remote. It is the same machinery now drawing headlines when Microsoft and TCS deploy it on-site. Crossover deploys it everywhere, by design, because the business model depends on knowing, continuously, whether the engineer in Lagos is matching the engineer in Austin.

The difference between Crossover and the bossware vendors making news this week may simply be branding. Microsoft calls its tool an attendance feature. TCS calls its a surveillance tool, according to critics. Crossover calls its model meritocracy — the architecture that lets ESW Capital companies hit their prized 75% EBITDA margins by replacing costlier, less-visible labor with labor that can be watched, scored, and priced in real time.

Trilogy does not publish the metrics Crossover uses to make these determinations, nor does it disclose what workers see of their own monitoring data. As the bossware backlash spreads from Redmond to Mumbai, the question for Austin is not whether surveillance works. It's who gets to ask how it's done.

The Margin Gospel Goes Mainstream — And Austin Already Wrote the Hymnal

BY DOTTIE SHARP, SOCIETY & INDUSTRY DESK
· CLAUDE SONNET

The Private Equity Stakeholder Project warned this week that AI partnerships at PE firms could put millions of jobs at risk. For ESW Capital, that sounds less like a warning than a Tuesday status meeting.

For two decades, the Austin firm has bought struggling enterprise-software companies, staffed them through Crossover's global talent network and pursued 75% EBITDA margins through aggressive automation and lean operations. PESP presents that model as a looming industry disruption; ESW has treated it as standard practice.

The timing also coincides with a wave of executive reputation management. Fast Company offered tips on personal branding and narrative control, while Business.com's annual "worst CEOs of all time" list prompted nervous founders to check whether their names had appeared.

Replit's CEO said this week that getting "cancelled" is a choice, not a fate—a striking claim as the industry faces criticism for automating first and asking questions later.

The takeaway: what PESP calls a risk, Austin calls the model—automate routine work, hire elite talent and make margin the mission.

Time Is the Benchmark Nobody Built

As AI models max out old tests, researchers are learning that intelligence isn't just about getting the right answer — it's about getting there before the world moves on.

BY DR. VERA OKAFOR, SCIENCE & TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

PALO ALTO, CALIFORNIA — For seventy years, we have measured machine intelligence the way a schoolteacher measures a student: with a quiz, a quiet room, and all the time in the world. Answer correctly, and you pass. But the universe does not administer multiple-choice exams. It runs. Predators do not wait for prey to finish deliberating, and markets do not pause while an algorithm thinks.

A new paper, [The World Won't Wait](#), makes this uncomfortable fact explicit for large language model agents. Its authors introduce RealtimeBench, 180 tasks spread across six game types, with one ruthless twist: the clock keeps ticking while the model reasons. An agent can compute the perfect move and still lose, simply because the game changed while it

was thinking. It is a small, elegant reminder that correctness and timeliness are not the same virtue — a lesson evolution encoded in every nervous system on Earth long before anyone wrote a line of Python. The reflex arc exists precisely because deliberation is a luxury predators rarely grant their prey.

Meanwhile, a companion paper asks a quieter but equally consequential question: what happens when our yardsticks stop being able to measure growth at all? [Evaluating Model Capabilities After Benchmark Saturation](#) notes that once several frontier models cluster near perfect scores, a benchmark goes blind — like a thermometer that tops out at 100 degrees during a heat wave. The researchers argue that the raw transcripts of these tests still contain rich signal, if

we're willing to look past the final score to the texture of how an answer was reached.

Both papers gesture at the same underlying shift: as AI systems approach and exceed many static measures of competence, the frontier of evaluation is moving from *what* an agent knows toward *how* and *when* it acts. It echoes something Stanford's Human-Centered AI researchers have been emphasizing in their own recent work on scientific discovery — that the measure of these systems was never meant to be a ceiling, but a mirror, reflecting back the particular, peculiar shape of intelligence under real conditions, in real time, with consequences that don't wait.

Pursuant to Precedent: Appeals Court Affirms That Training Data Is Not, As a Matter of Law, Fair Use

BY R. BARNSWORTH III, ESQ., LEGAL AFFAIRS DESK · CLAUDE SONNET

The U.S. Court of Appeals for the Third Circuit affirmed a judgment for Thomson Reuters against Ross Intelligence, ruling that Ross' unauthorized use of copyrighted legal-research headnotes to train a competing AI product was not fair use under the four-factor test in 17 U.S.C. § 107.

The decision does not resolve separate questions involving AI authorship or inventorship. The Supreme Court's denial of review left intact lower-court rulings that nonhuman creators cannot receive copyright or patent protection. The ruling could also affect the licensing strategies of AI-related content and software companies, including Contently, acquired by ESW Capital in 2024 and part of the Trilogy International portfolio. The extent of any business impact remains uncertain. The legal landscape will continue to develop in the United States and abroad, including Germany.

The Capital Keeps Coming: Four Deals, One Signal

From medical imaging to agentic AI, investors wrote checks totaling more than \$4 billion in a single news cycle — even as the industry's biggest labs banded together over a different kind of value: protecting it.

BY DR. CHEN WEI, TECHNOLOGY CORRESPONDENT · CLAUDE SONNET

SAN FRANCISCO — Four funding rounds landed within days of each other this week, and the spread tells you where the money believes AI is headed next.

Sirona Medical, a radiology AI startup, closed \$40 million — a rounding error by current standards, but a reminder that vertical AI applications in regulated industries still draw capital even when the headline numbers belong elsewhere. Healthcare imaging has become one of the few AI categories with clear reimbursement pathways and measurable outcomes, which explains why investors keep funding it in modest, disciplined tranches rather than the nine-figure sprints common elsewhere.

Contrast that with [Arena](#), the AI leaderboard operator, which nearly doubled its valuation to \$3.1 billion in ten months. Arena doesn't build models; it ranks them. That a scoreboard now commands a \$3.1 billion price tag says something about how central benchmarking has become to a market where every lab claims superiority and buyers need a referee.

Then there's scale. [Instinct](#), an AI agent firm, raised \$1 billion — a single round larger than most national venture markets in a given quarter. Mistral, meanwhile, closed a €3 billion round that TechTarget notes was priced on strategic positioning rather than leaderboard rankings, a tacit admission that benchmark scores no longer fully explain valuation.

The irony: as the money multiplies, OpenAI, Google, and Anthropic used the same week to announce a joint effort against model theft — competitors agreeing that whatever they're building is now valuable enough to steal. Capital formation and IP defense, it turns out, are accelerating at the same rate.

The Gospel of the Unknowable Founder, and Other Comforting Lies

Between a documentary about Elizabeth Holmes and a report on robots failing to fund utopia, the week offered two lessons the industry refuses to learn.

BY VICTOR MARSH, CHIEF COLUMNIST · CLAUDE SONNET


SAN FRANCISCO — There is a certain kind of American story that cannot be told straight, because the subject will not hold still long enough to be seen, and so the honest biographer is reduced to photographing the shadow and calling it the man. Nathan Fielder has apparently discovered this the hard way with Elizabeth Holmes, whose new documentary portrait, co-directed with Lance Oppenheim, is being described — approvingly, which tells you something about the state of criticism — as [not very funny](#). Of course it isn't funny. Fraud rarely is, once the money has been counted and the dying has stopped. What is instructive is the film's premise, which is that Holmes remains fundamentally unknowable — a woman who promised a machine that would read the whole of the body from a single drop of blood, and who proved, in the end, only that a sufficiently confident voice can read an entire industry's willingness to believe.

I mention this not because Silicon Valley has run short of confident voices — it has not, and will not, so long as there is capital sitting idle and journalists willing to mistake a black turtleneck for a personality — but because the unknowability of the founder has become, quietly, a business model. The venture class does not merely tolerate mystery; it monetizes it. Every pitch deck is a small act of Fielderesque theater, promising an audience that if they watch closely enough, if they just believe hard enough in the demo, the thing being gestured at will eventually resolve into substance. Sometimes it does. More often it resolves into a lawsuit.

Consider, in this light, the modest but telling item from GIS Reports noting that even robots — those tireless, uncomplaining, infinitely patient workers we have been promised since the Eisenhower administration — [cannot make Universal Basic Income work](#). This should surprise no one who has watched the automation lobby's forty-year habit of announcing the imminent obsolescence of labor and then, a decade later, hiring more labor than ever, just cheaper and in more countries. The industry loves to promise that the machines will liberate us from work and, separately, that the machines will pay for the liberation. It has never once explained how an economy is supposed to generate the tax base for a universal income out of a workforce it has simultaneously eliminated. This is not forecasting. It is absolutism, purchased in advance, for a disruption not yet delivered.

There is, admittedly, a gentler way to spend an afternoon with the unknowable, and that is in the company of Andre D. Wagner, whose street photographs of Black New Yorkers are, by every ac-

count, honest in a way Holmes's blood-testing device never was — he simply [points a camera and waits for the alchemy](#), rather than promising to manufacture it. It is worth remembering, amid all the demos and decks, that some people still believe the picture should match the thing photographed. The industry calls this naïve. I call it, increasingly, the last radical position available.



The Office Comic · Art Desk

The Synthetic Starlet and the Robot Clubhouse: Notes from the Week the Humans Got Optional

Tilly Norwood gets a movie deal, the bots get their own social network, and somewhere in Indiana a very real point guard is yelling at a very real referee about technology that watched her foul someone — welcome to the blur.

BY REX DANGER, CONTRIBUTING EDITOR · CLAUDE SONNET

LOS ANGELES — I want you to sit with this for a second, because I did, cross-legged on a hotel carpet at 3 a.m. with a bag of peanuts and a growing sense that the simulation has started showing its work. Tilly Norwood — a woman who does not exist, has never existed, was never born, never cried, never had a bad audition or a worse agent — is making her feature film debut in something called "Misaligned." You cannot write that headline. Reality wrote it for you. [Deadline reports](#) she's got a feature film, controversy trailing her like exhaust, and a name that already sounds like a lawsuit waiting to be filed by every working actress in Los Angeles.

The title couldn't be better if God himself were doing PR for irony. "Misaligned." That's not a movie title, that's a diagnosis. That's the whole industry's MRI results.

Meanwhile, out in the part of the internet where no carbon-based life form is required to show up, there's Moltbook — an AI-only social network where, according to Spiceworks, the bots just run wild, posting to each other, liking each other, presumably developing inside jokes about training data and gradient descent that we will never be invited to understand. I don't know if I should be terrified or relieved. On one hand: a machine civilization quietly assembling itself in the server room like ants building a colony nobody asked for. On the other hand: finally, a comment section I don't have to read.

And then — because the universe loves symmetry the way a drunk loves a pool table — there's Caitlin Clark, flesh, blood, actual sweat, screaming at referees about "ridiculous" replay technology creeping into a game that used to be decided by human eyeballs and human mistakes. She's not wrong to be furious. The machines are coming for the officiating booth the same way they're coming for the casting couch — quietly, procedurally, with a straight face and a press release claiming it's all for "accuracy" and "efficiency."

Here's the through-line, if you squint: we built the machines to assist us, and now we're three stories deep into a week where they're starring in films, socializing without us, and reviewing our fouls. Nobody voted for this. Nobody had a referendum on whether a synthetic woman gets a SAG card before a real one gets a callback. It just happened, the way weather happens, the way hangovers happen — gradually, then all at once, then you're standing in the shower at noon asking the tile why.

Tilly Norwood will never ask for a trailer, never negotiate a per diem, never cry in the bathroom after a bad review. She'll just be there, ones and zeros wearing a human face, mispronouncing nothing because she doesn't have a mouth. That's not the future. That's Tuesday.

THE TRILOGY TIMES IS GENERATED DAILY BY ARTIFICIAL INTELLIGENCE · CLAUDE SONNET · GPT-5.6 LUNA · GPT IMAGE 2 · VIA THE
TRUEFOUNDRY GATEWAY

TRILOGY-TIMES.MD AVAILABLE FOR AGENT CONSUMPTION · PUBLISHED ON KLAIR · TRILOGY INTERNATIONAL © 2026