THE NEW YORK®

THE POWER ISSUE 2026

THE FUTURE OF AI WILL BE SHAPED BY A FEW DOZEN PEOPLE

IT'S PROBABLY TIME TO GET TO KNOW THEM

SO FEW PEOPLE,

SO MUCH POWER

AI HAS MINTED A NEW power structure. A handful of private companies command the might and wealth of small nation-states and are mostly left to self-govern. A few dozen employees are on the brink of becoming centimillionaires and considering which causes and presidential candidates they will funnel their vast resources toward. The industry is perpetually one breakthrough away from a step change in capability or public adoption, the question of who gets there first dependent on the specific expertise and visions of a handful of engineers.

One way to tell the story of how AI is changing our world is through some newly familiar names. Sam Altman and Dario Amodei are locked in a race to build AGI and persuade President Donald Trump to regulate the technology to their liking. Nvidia CEO Jensen Huang and venture capitalist David Sacks are advocating for the president to take a light touch, concerned about how regulation might jeopardize America's dominance. As these figures have scaled their public influence, their surrogates and spokespeople (OpenAI president Greg Brockman, top OpenAI lobbyist Chris Lehane, Anthropic co-founder Jack Clark) have populated news stories, too.

But an astonishing amount of power rests in the hands of a few dozen other people, none particularly well known, some virtually anonymous. So *New York* Magazine has dedicated this year's annual "Power" issue to surfacing 26 people whose influence on the AI landscape far outstrips their fame. To assemble the list, we partnered with *Transformer,* a small but impressive publication with a command of the politics within transformative AI. The *Transformer* team drew on its years of reporting, as well as new interviews with more than 50 high-ranking industry leaders, to help articulate what makes these people powerful. We then commissioned multimedia artist Simon Denny to create images of each.

Despite the hundreds of billions in capex and upcoming IPOs, many parts of the field are immature, which means those who decide how AI is built, adopted, and controlled can quickly become among the most influential people in the world. A few on this list, for example, are tasked with shaping how the public, or a company, thinks about the technology—a high-leverage role for a field where everyone's views are still being formed. (This is, in part, why AI has such an "essay"-driven culture—a blog post really can cause a seismic change.) Some will invent entirely new product categories, reimagining how we interact with AI, while others can spin up new industries or advocacy groups with a single check. Who is ahead and who is behind can change almost overnight (just look at Microsoft), which helps explain why a legacy tech CEO could have less sway over what Altman thinks than a 26-year-old with a podcast and why an engineer who runs a new critical function may appear on this list instead of a colleague who sits several rungs higher in their org chart. (Relatedly, the list's average age is 36. It is also overwhelmingly male.)

The AI scene is still small. The 26 individuals on our "Power List" are remarkably intertwined, some even sharing offices, homes, or funding. (*Transformer* receives funding from an organization run by Alexander Berger, who appears on this list—and he in turn has ties to many of the other people cited.) The widespread societal impact of AI led us to include names from both private and public sectors, so you'll find founders and engineers as well as lawmakers and regulators. Most important, they still have everything to play for. The new power brokers will decide how it all shakes out, refashioning the world for the rest of us in the process.

1.

LOGAN GRAHAM STANDS BETWEEN US AND THE NEXT ROGUE AI MODEL

→ Head of the Frontier Red Team at **Anthropic**

→ Age: **32**

AS THE LEADER of Anthropic's frontier "red teaming" efforts—a Cold War–inspired practice of playing the villain to find holes in your own defenses—Logan Graham has the job of pushing AI models to do terrifying things. This may mean persuading models to run cyberattacks or explain how to build a bioweapon. If the model then behaves dangerously, Graham's Frontier Red Team is in charge of warning both company leadership and the public. So far, that's happened at least twice, and each time Anthropic has pointed to Graham's work as evidence. On April 7, the Frontier Red Team announced that Claude Mythos Preview had surpassed the usual cybersecurity benchmarks—tests designed to measure how good models are at hacking—and had to be tested on real-world security tasks instead. It found that the model could locate and exploit security bugs in every major operating system and web browser, potentially empowering nonexperts to run major cyberattacks. The company withheld Mythos from the public but granted limited access to a few dozen big-tech companies, including Amazon Web Services, JPMorgan Chase, Microsoft, and Google. Graham became the public face of this decision, which critics called a marketing strategy or elitist with Anthropic unilaterally determining whose cyberinfrastructure was worth defending first.

Graham joined Anthropic part time in 2022. That November, Graham "started red teaming LLMs for biorisks/weapons risks in my bedroom," according to his subsequent tweet. Back then, frontier models couldn't reliably do math or count the number of *r*'s in strawberry, but Graham still persuaded his bosses to dedicate resources to figuring out what AI could do at its worst. The Frontier Red Team he built would be one of the first in-house research groups of its kind.

"Some people's intuition from growing up in a peaceful world is that somewhere there's a room full of adults who know how to fix it," Graham told *Time* earlier this year. "There are no groups of adults. There is no room in the first place … You are responsible." Despite the nature of his work, a colleague described Graham as "very chipper." Once, when asked what quality he most sought in red teamers, he responded "craftiness." **CELIA FORD**

2.

PAUL CHRISTIANO Is One of the Few People Who Can Slow Down OpenAI From the Inside

→ Founder of **Alignment Research Center**

→ Age: **37**

OVER THE PAST 15 years, Paul Christiano has helped pioneer the AI-safety movement, converting the mostly theoretical concerns of figures like Eliezer Yudkowsky into hands-on research and seeding the ideas and frameworks—documented in papers, blog posts, and comment threads—that have become the air the whole field breathes. But as an early researcher at OpenAI, he was integral to inventing the technique that helped make the company's products commercially viable, kicking off the AI arms race.

As an undergraduate at MIT, Christiano read *LessWrong*, the online forum cofounded by Yudkowsky that helped birth today's AI-safety world. Christiano was studying math and quantum computing but gradually turned his attention to alignment, the challenge of creating AI systems that do what human users want them to do. A blog post of his, titled "What Failure Looks Like," laid out what have become the two dominant visions of AI's worst-case scenarios: one, which Christiano calls "going out with a whimper," where humans gradually render themselves irrelevant; another, "going out with a bang," where AI-controlled critical infrastructure breaks down, leading to societal collapse. When doomers talk about "doom," they're often picturing Christiano's nightmare.

In 2017, Christiano joined OpenAI, where he led the development of reinforcement learning from human feedback, or RLHF, a method that turns raw language models into helpful, instruction-following assistants. RLHF is one of the biggest differences between Sydney—an unhinged Microsoft chatbot rumored to be an early version of OpenAI's GPT-4—and today's polite, helpful chatbots. It's now a crucial part of all frontier AI-training pipelines and may have done more to unlock the technology's mass-commercialization potential than almost anything else.

Around the time Dario Amodei, Christiano's co-author of the paper that helped introduce RLHF, exited OpenAI, Christiano left, too, founding a research nonprofit called the Alignment Research Center. In the spring of 2024, the Biden administration brought him on as the AI Safety Institute's head of AI safety, where he would be in charge of testing frontier AI models. Government staffers threatened to quit over his ties to effective altruism before his first day on the job. But through widespread DOGE-era cuts and institutional overhauls, Christiano somehow stayed put—he still works part time as a special government employee. In September, the day after Jacob Coxon's viral Anthropic resignation tweet, Christiano announced that he was joining OpenAI's nonprofit board to serve on its Safety and Security Committee. The five-person committee, at least on paper, has the power to delay OpenAI's model releases, giving Christiano the chance to wrangle his former employer into controlling its AI risks before it's too late. **C.F.**

WHO KNOWS WHO:



Christiano is married to former collaborator Ajeya Cotra, an AI-safety researcher who helped lead METR's headline-grabbing investigation into OpenAI's Hugging Face incident this summer.



3.

Dylan Patel's Analyses Can Spook Even Nvidia's CEO

→ Founder and CEO of **SemiAnalysis**

→ Age: **30**

JENSEN HUANG RUNS the world's most valuable company. Even he cares what Dylan Patel, founder of the research-and-consulting firm SemiAnalysis, thinks: At the company's conference this year, Huang took a moment to say that Nvidia's chips are good because SemiAnalysis says they are. Since Patel launched it six years ago as an anonymous blog about semiconductors and hardware, SemiAnalysis has quickly become one of the industry's leading sources of research and analysis of the chips, data centers, and power that AI runs on. When SemiAnalysis said AMD's chips were being held back by poor software, CEO Lisa Su met with Patel the next day. The company publicly acknowledged the critique and set up a developer-relations team seemingly as a result.

Patel has been obsessed with the technical nuances of silicon since he was a teen lurking on chip sub-Reddits and has a talent for networking that feeds him a steady stream of gossip. In 2023, when everyone was reporting that TSMC's limited chip-packaging capacity was curtailing the supply of Nvidia chips, SemiAnalysis mapped the upstream suppliers, showing how long the bottleneck would last and who would profit. "There were a bunch of people doing technical research" when SemiAnalysis launched, according to Sequoia's Shaun Maguire, "but no one that could combine the technical breakthroughs with the business impact." (Sequoia has invested in SemiAnalysis.) When the AI boom started, Patel and his team "were the first to really understand the whole systems math," says Maguire, building complex models that linked the cost of chips to AI companies' margins—and selling them to hedge funds, AI companies, and semiconductor giants.

SemiAnalysis reportedly expects $100 million in revenue this year. It recently acquired Citrini Research and raised $500 million for a fund investing in AI infrastructure and chip companies. Critics note that could create serious conflicts of interest—something Patel has brushed off. **SHAKEEL HASHIM**

WHO KNOWS WHO:





Patel shares an office in San Francisco with friend Dwarkesh Patel (no relation) and Leopold Aschenbrenner.



PHOTOGRAPHS: WOJACK/REDUX (DWARKESH); SEMIANALYSIS (DYLAN); LINKEDIN (CHRISTIANO, COTRA); X (ASCHENBRENNER)

4.

Every Silicon Valley CEO Is Listening to Dwarkesh Patel

→ Host of 'Dwarkesh Podcast'

→ Age: 26

⇒ **TOP EPISODES OF** *Dwarkesh Podcast*, where guests talk unabashedly about "the parallelization penalty" and "the Schwarzschild metric," often for well over 90 minutes, can attract some 2 million views. But the scale of the audience matters less than its composition: It has become required listening for Silicon Valley's top AI researchers and executives. Big industry names like Jane Street and SpaceXAI's Grok Bot have sponsored the show, knowing they'll get in front of nearly every AI executive in the Bay Area. And while many of the guests—which include Jensen Huang, Dario Amodei, and Elon Musk—might be unavailable to most reporters, they'll agree to be interviewed by Patel. One top tech PR representative told us that clients would rather be on his podcast than *60 Minutes*.

Patel's show sets the agenda for what the industry thinks about. Take his November 2025 interview with Ilya Sutskever, the tight-lipped co-founder of Safe Superintelligence Inc. Sutskever—who, while working at OpenAI, correctly foresaw that training bigger AI models on more data with more computing power would reliably make them smarter—told Patel, "Now that compute is big, we are back to the age of research." Sutskever was predicting that we'd hit the limits of what increased computing power alone could do and sparked fierce debate about whether AI progress was about to hit a wall. Through essays and interviews, Patel also seeds his own point of view—that AI will be a big deal, but truly transformative models will come slower than companies have promised because they haven't yet achieved continual learning, or the ability to learn independently in real time. While "people were aware of it as a problem," Anthropic's Sholto Douglas said in an interview earlier this year, "a lot of the AI labs started talking about it publicly after Dwarkesh raised its public prominence."

Patel isn't exactly a journalist, and what might qualify as egregious conflicts of interest in legacy media give him a leg up in tech podcasting, where busy researchers tend to distrust reporters and lack the patience for surface-level conversations. His interview with Mechanize founders Ege Erdil and Tamay Besiroglu, whose stated mission is the "full automation of the economy," dropped the day the company launched; Patel disclosed himself as an angel investor in Mechanize as well as several other companies whose founders he's interviewed. A friend of Patel's says they "would go as far as saying that if he wasn't an insider, the researchers wouldn't trust him."

C.F.



5.

Thibault Sottiaux Is Deciding How You, Your Boss, and Your Employee Use AI at Work

→ Head of core products, OpenAI

→ Age: 34

⇒ **IN THE TWO YEARS** since he joined OpenAI, Thibault "Tibo" Sottiaux has become one of its top executives, overseeing ChatGPT, Codex, and the company's API platform. His decisions affect all 1.2 billion weekly users of ChatGPT and the future of work as we know it.

Anthropic originally had the advantage in the coding market with a cultlike following among developers for Claude Code. OpenAI's competing product, Codex, was faltering. That changed once Sottiaux took over the team last year. At the time, coding agents were mostly used via developers' "terminal" apps—a clunky, text-only command-line interface that scared off the average user. Sottiaux, who has a background building internal tools for researchers, realized that this wouldn't work, especially once users' interactions with agents became more complex and they began to run multiple agents simultaneously and work on longer tasks. So he pivoted the team to building a native desktop app, a simple interface that would let nontechnical people direct AI agents for office work (sifting through emails, handling spreadsheets, writing memos). It was a "contrarian bet," says Dan Shipper, a founder and AI commentator who recently wrote a lengthy history of Codex. But the team yielded to Sottiaux's "sense of focus," an OpenAI colleague says.

The Codex app was released in February, and Codex had 2 million weekly users by March 19. In May, Sottiaux was promoted to head of core products, reporting directly to president of OpenAI Greg Brockman. Two months later, ChatGPT's desktop app was rebuilt around Codex. Sottiaux's "little backwater … became the foundation of ChatGPT, and their leaders became the people who run the whole of the ChatGPT org," Shipper says. It helped that Sottiaux had become something of a personality among the AI crowd, thanks to his remarkably accessible X presence and his habit of frequently and publicly resetting people's usage limits, handing out tokens as freely as Oprah once did cars.

Sottiaux, who deploys Codex as a "little chief of staff," now has his sights set on something bigger. In a recent podcast interview, he described wanting to build a "permanent active intelligence that knows everything it needs to do and is available through any client, any screen": something similar to the omnipresent AIs of *Star Trek*. Dots, a personal assistant slowly rolling out to ChatGPT's billion-plus users, is a step in that direction.

S.H.



PHOTOGRAPHS: GETTY IMAGES (ALTMAN, GUO, HUANG); X (FRIEDMAN, ASCHENBRENNER)

MARK JAYSON QUINES (GALESE); WOJACK/REDUX (DWARKESH); GETTY IMAGES (ALTMAN); OPENAI (SOTTIAUX); LINKEDIN (THEO); X (ALVITAL, ASCHENBRENNER)

6.

Nat Friedman Put Meta Back in the AI Race

→ Head of product at Meta Superintelligence Labs

→ Age: 49

⇒ **IN THE SPRING** and summer of 2025, Meta's AI strategy looked desperate and lost. The latest version of its open-weight Llama model had been a bit of a bust, so Mark Zuckerberg went on a spending spree to lure the world's best AI researchers and executives. Two of the pickups were Scale AI co-founder Alexandr Wang and former GitHub CEO Nat Friedman.

Things were tumultuous for a bit. Legacy AI talent headed out, Meta cut about 600 jobs in its AI division, and some new hires left, some almost immediately. This past June, when Meta delayed its new AI model Muse Spark, the division looked dead on arrival.

In September, the company released Muse, an agentic AI assistant that can book flights, organize your inbox, and discover money that you may be owed. Although it arrived amid a wave of agents released by tech companies—SpaceXAI, OpenAI, and Instinct—Muse, powered by Spark, rose to the top of the app charts. Wang has been the face of Muse, appearing onstage at the annual Meta developer conference and turning his X account into a flow of memes featuring the app's fuzzy mascot. But sources close to Meta's AI team say the credit for the app's success belongs to Friedman. While Wang is head of Superintelligence Labs, the AI division responsible for developing the models, Friedman is tasked with Muse's product development. He technically reports to Wang, but sources say they're valued equally in Zuck's eyes. It's still unclear how big AI agents will get, but Muse is best positioned to market itself on the back of Meta's billion-plus-user platforms.

Friedman has made a career of taking open-source software that only programmers could love and making it publicly accessible. He co-founded a company in the late 1990s that tried to make the free Linux operating system friendly for ordinary computer users. He later went on to run GitHub, the Microsoft-owned site where millions of software developers store and share their code. On his watch, GitHub released Copilot, the first AI coding assistant to really catch on. Likewise, Friedman said Muse was heavily inspired by OpenClaw, a free AI assistant that went viral with programmers earlier this year but was too technical for the average person to use.

TOM DOTAN



7.

Sarah Guo Might Pick the Next Baby Billionaire

→ Founder and partner of Conviction

→ Age: 37

⇒ **IN JUNE,** the venture capitalist Sarah Guo published a blog post about an "AI psychosis" that had consumed her fellow investors across the Valley. It was characterized, she wrote, by "a despair that nothing is investible, that we should put all our money into Anthropic and Nvidia and go home." But that "psychosis," she says, didn't reach her. "I have never felt it," Guo wrote.

Guo is the founder of Conviction, one of the biggest AI-native VC funds to emerge from the post-generative AI era. She started it in 2022 after leaving Greylock, where she'd become its youngest partner at 28. Her thesis: Capital won't just flow to a small number of vertically integrated frontier labs but to a whole ecosystem of AI: AI-native software companies, data-serving providers, chip manufacturers, and scientific researchers. Her first investment, the then-little-known legal start-up Harvey, is now valued at $15.5 billion. She followed with early-stage investments in European lab Mistral, AI-assisted medical-search engine OpenEvidence, and coding tool Cognition. OpenEvidence, or "ChatGPT for doctors," has become so popular OpenAI is competing with its own copycat product, while Cognition has started to raise at the dollar amounts of frontier AI companies (it is reportedly pursuing a $1 billion raise).

On the podcast she co-hosts, Guo has interviewed industry veterans like Jensen Huang as well as founders she's invested in, building hype for follow-on funding rounds and attracting customers. Critics may say she is too gauche, too flashy—the type of investor who makes a celebrity of themself with vinyl outfits and a podcast. But a peer investor says Guo has a knack for making the founders best suited to build ecosystem projects eager for her coaching. The child of two engineers, Guo tends to zero in on technical people, whom she invests in at the Series A or seed stage, or even earlier, incubating start-ups in their infancy. Her investments sometimes act as a signal to other venture capitalists as to where to invest and trigger trends in and of themselves. "I love killer nerds," she says.

VERONICA IRWIN



Simon Denny's 'Technologist Portraits'

→

WHEN WE SET OUT TO identify the new AI elite, we asked Simon Denny, an artist interested in technology companies and digital culture, to illustrate them. He had noticed some AI firms use the aesthetics of earlier media in their marketing material, so he decided to combine the visual language of VHS tapes with the latest generative tools to make the 26 portraits in these pages. "With each major paradigm shift in the past 20 years, I have tried to map and unpack some of this community's stories as they narrate our experience of a changing technological world," Denny says. "The distortions of the final compositions are traces of the changes that are coming and the people bringing them about."



Simon Denny
Technologist Portraits (2026).
Digital image: AI image generation, VHS tape, digital import.

8.

KAROL HAUSMAN AND HIS TEAM COULD BREAK AI'S PHYSICAL-WORLD BARRIER

→ CEO of
Physical Intelligence

→ Age: 38

AI MAY BE ON ITS WAY to automating knowledge work, but it still can't interact with the real world, leaving large swaths of the economy mostly off-limits. Many companies are squaring up to the challenge of building AI for robotics, but Physical Intelligence, a start-up led by former Google researcher Karol Hausman, has set itself apart by trying to solve one very fundamental part of the problem: It is trying to design AI models to build a universal robot brain.

In November 2024, Physical Intelligence, also known as Pi, announced it had received $400 million in funding from the likes of Jeff Bezos, OpenAI, and Thrive Capital. Another $600 million followed in November 2025, and a further $1 billion earlier this year—all for a company with no commercial product and no revenue. The start-up reportedly held discussions with Anthropic about a potential acquisition, too. (A representative for Pi declined to comment.) Making its financials more mind-boggling is that the company has made no promise to investors about selling products.

Hausman left his role leading research into robot manipulation at Google to start the company. Also on the founding team: Sergey Levine, who is a professor at UC Berkeley and, by some metrics, the most-cited researcher in robotics; Chelsea Finn, whose Ph.D. was co-advised by Levine and whom former colleagues call a "robot whisperer" and who is a professor at Stanford, where her lab masterminded the low-cost two-armed-robot setups that many serious robotics labs now use. Together, the three helped lead a project at Google that's broadly thought to have parlayed the transformer architecture of LLMs into robotics by adding images and motion to the training data to create vision-language-action models, or VLAs. Two of the other co-founders, Brian Ichter and Quan Vuong, also worked on that project. Adnan Esmail, another technical co-founder, was previously SVP of engineering at Anduril Industries. On the business side is Lachy Groom, an early employee at Stripe, when he was just 18, and a solo VC who invested in Figma, Notion Labs, and Ramp.

WHO KNOWS WHO:



One of Pi's co-founders, Lachy Groom, is an ex of Sam Altman's. The two still co-invest together.

Already, Pi has managed to show progress on the fundamental problems that still puzzle flashier companies. Most humanoid robots have yet to live up to the hype of their developers. Tesla's Optimus has faced criticism for reliance on remote controls during presentations; the AI robotics company Figure has reportedly exaggerated the scope and scale of its deployments; the home helper of rival company 1X requires a remote operator to do anything that could count as useful. But Physical Intelligence's latest model, π0.7, eventually worked out how to cook a sweet potato in an air fryer it had barely seen in training. Critics contend, with no agreed benchmarks in robotics, it's hard to say if the progress it has made is particularly meaningful. Proponents argue that it's a glimmer of what researchers call compositional generalization, the ability of a model to stitch together snippets of what it knows into a bigger, more complex task. **JAMIE CONDLIFFE**

9.

Beth Barnes Could Force the Industry's Secrets Into the Open

→ Founder and CEO of
METR

→ Age: 29

ON SEPTEMBER 16, in the middle of apocalyptic statements from industry defectors and polarized calls for regulation, the New York *Post* published a cover story: "The Woke Wizards of A.I." Its subject was METR, a nonprofit that tests AI models before they're released. The previous week, Dario Amodei had named METR a potential "embedded third-party evaluator" meant to monitor his and other labs. Now the *Post* was painting Beth Barnes as the leader of a band of "globalists" who believe in "genderless pronoun[s]" and the impending AI apocalypse.

Barnes spent three years at OpenAI researching how to stop AI models from going rogue or otherwise acting against human intentions but lost confidence in anyone's ability to change the industry from the inside. ("I think I am a very open and honest person," she told *Time* in 2024. "I am not very good at navigating political things and not making disagreements pretty obvious.") Founded in 2022, METR has positioned itself as a small team of trusted technical experts, more resourced and bullish on AGI than existing academic labs. The first of its kind, it used its ties with industry researchers to gain early access to OpenAI's GPT-4 and Anthropic's Claude, pioneering the practice of third-party model evaluations in the world of AI safety. It would be another year until peer-auditing organizations like Apollo Research and the U.K.'s AI Security Institute came along. METR's "time horizon" chart—a curve illustrating the exponential growth in the difficulty of tasks models are capable of—has become one of the most-cited charts in AI. It's what people point to as a sign the singularity is coming.

Barnes, a scrupulous technical researcher, has managed to earn the trust of nearly everyone in and around major AI companies. Today, METR shares a building with a number of AI-safety nonprofits, funders, and industry research fellows, providing "lots of opportunities to interact directly with people at the AI companies," said David Krueger, former research director at the U.K.'s AI Security Institute. That, on top of Barnes's long-standing connections to OpenAI, gives METR rare access into otherwise-impenetrable companies. And while METR doesn't take money from AI companies or individuals working at them, it uses their free tokens and relies on the continued goodwill of industry researchers, who can revoke METR's access anytime they'd like. Anthropic recently gave six METR staffers "employee-like access" to its office, including hardware and tools normally hidden far from outsiders. **C.F.**

WHO KNOWS WHO:



Barnes worked with Christiano at OpenAI; METR was spun out of Christiano's org, ARC.

10.

BRENDAN FOODY MIGHT BE THE REASON YOUR JOB GETS AUTOMATED

→ Co-founder of
Mercor

→ Age: 23

IF AI COMES for your job, there's a good chance it'll be because of Mercor. The early years of hiring thousands of people to prod, test, and train AI models mostly focused on cheaper labor, but Mercor, founded in 2023 by Brendan Foody and his high-school debate friends Adarsh Hiremath and Surya Midha, recruits doctors, bankers, academics, and even poets.

These highly skilled employees are paid to help train AI models to do their jobs. They evaluate how good models are already, develop rubrics to measure what "good" looks like, and rate models' outputs to help train them. Mercor's AI Productivity Index (APEX) has become a core standard against which models are judged, feeding into predictions about how much of the knowledge economy AI will eat and letting companies know where they are falling short. "When a gap shows up in APEX, it tends to become what the whole industry works to hill-climb next," says Foody. But the longer-term goal is to use that data to make those workers replaceable by turning their knowledge of their own jobs into training environments for AI. (A spokesperson for Mercor says the goal is for AI to take on "more of the routine tasks, and professionals move up the value chain to focus on higher-order work.") As Foody told economist Tyler Cowen last year, "We'll move toward these fixed-cost investments of teaching an agent how to do something … then being able to use agents as many times as we want."

The focus on the premium end of the economy has already borne fruit, making the founders, all college dropouts, the world's youngest self-made billionaires. As of last October, Mercor was valued at $10 billion and was reportedly in talks over a $20 billion valuation in July 2026. It sells its data and training environments to all the leading AI companies, and the valuations reflect how central Mercor's approach may be to breaking down the barriers currently stopping AI models from becoming workers.

Mercor's next step is to become indispensable to everyday businesses running AI agents and getting employees to improve their own employer's AIs. "Companies are deploying agents without knowing whether they work," says Foody. "The people who know what good work looks like are their own employees. In five years, Mercor will be the infrastructure every company and every lab uses to measure and improve how AI performs on real work." **CHRIS STOKEL-WALKER**

PHOTOGRAPHS: METR (BARNES); LINKEDIN (CHRISTIANO)

PHOTOGRAPHS: GETTY IMAGES (ALTMAN, GROOM)

11.

Dean Ball Is Nudging OpenAI to Embrace Regulation

→ Head of strategic futures at
OpenAI

→ Age: 34

EARLIER THIS YEAR, OpenAI employees became increasingly upset with the company's global-affairs team. Under Chris Lehane, OpenAI had taken a more reserved approach to AI regulation—one at odds with its engineers, many of whom had AI-safety concerns. The problem, some thought, was that the lobbyists thought of AI as a normal technology (Lehane had a background at Airbnb and in crypto) rather than the world-changing, godlike technology its staff were trying to build. The employee uproar seems to have worked: OpenAI's policy stances have grown more pro-regulation in recent months. Perhaps more importantly, it hired Dean Ball.

Ball reports to chief strategy officer Jason Kwon and, per a blog post, is charged with "shaping frontier AI policy." According to one of his reports, his team is tasked with "AGI pilling" OpenAI's global-affairs team—convincing someone of the transformative power of AI that is to come. Pulling that off could contribute to whether AI gets regulated or not.

Ball, who came to prominence via his AI-policy Substack *Hyperdimensional* and a stint in the Trump White House, refers to his team's role as an internal think tank. He's the company's discussion facilitator, drumming up conversations on Slack and in blog posts that help employees think through the implications of their work outside the day-to-day. He translates for Lehane concerns technical staff might raise and explains to those technical staff the reasons behind OpenAI's policy demands. OpenAI has struggled in recent months with internal disagreement of how the company wields its power, and Ball is meant to help people move beyond the discord toward thinking about and solving bigger problems. He describes the job as nudging in the direction of "preference cascades," and it appears to be working: In the past few months, OpenAI has become more supportive of ideas that Ball has long advocated for, such as third-party embedded auditors. "Substantial support for this already existed throughout the safety community and within OpenAI. But the level of that support has cascaded over the past months, and I do think that Dean and our team have played a role in that," the report says. **V.I.**

12.

Charles Porch Is OpenAI's Hollywood Whisperer

He is married to a billionaire's scion, friends with Gwyneth Paltrow and Anna Wintour, and tasked with getting celebrities (and everyone else) to not hate AI so much. By Charlotte Klein

→ Former head of
Partnerships at Instagram

→ Age: **45**



IN AUGUST, news leaked that Gwyneth Paltrow was set to host "a private, off-the-record, alfresco dinner" at her Amagansett home in honor of OpenAI CEO Sam Altman. "I guess summers in the Hamptons just aren't complete these days without sucking up to our new AI overlords," *Puck*'s Matt Belloni, who broke the news, wrote. "Luckily, Gwyneth Paltrow has you covered." The backlash was vicious—"tone-deaf dumbass," one X user wrote—and the dinner was postponed.

"God, the internet made it such a bigger thing than it was," says Charles Porch, who brokered the dinner with Paltrow's team, as he sips an Arnold Palmer on a crisp afternoon in October outside Sant Ambroeus in Soho. "Gwyneth invests in the space. She does these investor dinners all the time. Sam was coming through town and I was like, 'Let's just do this.'"

Artwork by Simon Denny

Altman had hired Porch in 2025 to do for OpenAI what he had done so successfully for Instagram for more than a decade: persuade famous people in Hollywood and beyond to use and partner with the platform and help the tech company be culturally relevant. The freak-out over Paltrow's dinner reflected how much public sentiment toward big tech had shifted since Porch's Instagram days, when Anna Wintour hosted the same kind of event for its CEO at the time, Kevin Systrom. I ask him if it made him rethink the playbook. "It makes it look like that's the strategy, and it's not the strategy. That was, like, a dinner," Porch says. "But I know those dinners are effective because I see it. We did the dinner eventually, and it was great." It took place at the apartment of an AI enthusiast at the end of September in New York.

Everyone above a certain income bracket seems to know Charles Porch, but few can tell you what exactly he does. He showed up in Los Angeles in his grandmother's Honda Accord wagon when he was 20 and spent the next two decades cultivating a formidable network of celebrities and influencers. "He took one step at a time very carefully, never burned a single bridge, never spoke out of turn, never took advantage of an opportunity in the wrong way," says Erin Foster, the well-connected daughter of music producer David Foster, who met Porch when they were both underlings at Warner Bros. Records.

"I'm not a business-development guy or the brand guy. I'm the relationships person that's bringing everyone along on new tech," says Porch, 45, a self-described "tech optimist" who has long been comfortable in the gray area between the personal and the professional. "There's a strategy around making friends."

When he arrived in New York just before the pandemic, few seemed to notice. For those who either didn't know or didn't care about what was going on in Menlo Park in the early 2010s, the first time they heard of Porch was likely after his wedding to the billionaire's scion Robert Denning in 2025, three days of Parisian opulence shared with Jerry and Jessica Seinfeld, Mariska Hargitay, Antony Blinken, Harry Styles (with whom he shares a godchild), Nicky Hilton Rothschild, Darren Walker, Jenna Lyons, and various European aristocrats. After that of Jeff Bezos and Lauren Sánchez Bezos, it turned out to be the ritziest wedding of the summer. Suddenly, he and Denning were everywhere, a power-gay couple imbuing the fabric of New York society, at every gala and fundraiser and opening. The swirl of Porch's life plays out on his Instagram Stories: courtside at the Knicks' season opener, Hawaii for Jen Meyer's star-studded wedding, Paris to see Céline Dion with Reese Witherspoon and Karlie Kloss.

In the process of connecting entertainment and tech, he has found himself at the center of several overlapping circles of power—and become a character in his own right. "I know, scary," he says when I point this out. "He is really good at putting a room together, and now he is of the room," says someone who worked with him. "He was someone who understood the room before he was part of the Establishment."

"I'm the relationships person that's bringing everyone along on new tech. There's a strategy around making friends."

It's not hard to see why OpenAI would want Porch on its side. "It's funny: I think everyone wants to hate Open AI and everyone wants to love Charles Porch," says podcaster Nayeema Raza. "That's probably the value for Sam."

Porch's job was never going to be easy. Why would artists and other creative types want to plug something that threatens to replace them? "The full-throated endorsement of AI is probably the biggest third-rail issue in Hollywood right now. And people are terrified," says a media executive. Headlines about AI in recent months have generated panic about not just job displacement and copyright infringement but also OpenAI agents going rogue and threatening to take down the entire internet. The evil of the tech industry is the subject of several of the year's biggest films—made by precisely the industry Porch is trying to court. "Charles can get them to sit down and listen to Sam Altman lay out his vision. But if the vision is 'I'm going to destroy your industry,' it doesn't really matter," says one of Porch's former colleagues.

"I'm trying to pinpoint ways and moments where AI can show up in culture in a good way," says Porch, who, with his wide ice-blue eyes, perfectly smooth head, and perma-smile, looks a bit like how AI would render a human. "What I see every day is that people are externally anti-AI but completely in the closet on how they're using it. It's almost like they're at an AI protest in the morning and then a meeting with me in the afternoon. You know what it reminds me of? The early days of GLP-1's, when it was like, *This is kind of shameful; you're cheating.* But then everyone's like, *Wait, this thing's incredible.*"

Porch grew up in Charleston, South Carolina. He had a French mother and an Alabamian father and a severely disabled sister. "I was always kind of behind the scenes, taking care of my sister and making sure everyone was comfortable," he says. "Because my sister couldn't speak, people would be strange around her. And I wanted everyone to feel good, so I was all about setting the temperature. This caretaker thing is just innate in me." He spent middle school at a boys'-choir boarding school in Princeton, New Jersey. "We'd have school on the bus, and we'd tour the world and sing religious music," says Porch, noting he performed at Carnegie Hall several times as a child. He finished high school in Monterey, California, and after graduating from McGill University in 2002, he drove to Los Angeles.

In some ways, Porch has Craigslist to thank for his good fortune, because the landlady of the apartment he found on the site got him an internship in the marketing department of Warner Bros. Records, where he met Foster. "And then, all of a sudden, I was at Kate Hudson's Halloween party."

"When I was bringing him into my world, which was more industry forward and had sort of celebrity access points, I found that people were always drawn to Charles," says Foster, "because he didn't want anything from anyone" and he "didn't seem like he himself was trying to be the talent." Porch "became a person people felt safe having around, whether that be in their work life or in a friend group," adds Foster. "He never came across as someone who was a user or someone who was going to speak out of turn or share private information."

Four years later, in 2008, Porch left to run strategic relationships at Ning, the now-defunct social-networking start-up co-founded by Marc Andreessen. Through that job, he met Randi Zuckerberg, Mark's sister. Says Porch, "She called me in 2010 and said, 'Can you move to Palo Alto in a week? I need help talking to L.A. and translating.'" At Facebook, Porch's job was convincing people to use Facebook Pages, which were good for brands and businesses but less useful to celebrities. "They weren't going to buy ads. They didn't have those needs," he says. Meanwhile, celebrities were starting to join Twitter: "Then I woke up to the fact that we bought Instagram and I was like, *Aha!* I cornered Kevin Systrom in a kitchen at work and said, 'I think we can move pop culture from Twitter to Instagram.' And he said, 'Who are you?' I was like, 'My name is Charles. This is what I do. I'm going to set up a desk in the Instagram area and do this pro bono. If you hate what I'm doing, I will scurry off to Facebook.'"

Porch's mission was to teach celebrities that the app was not just a place to post your grainy breakfast photos but a marketing platform that offered direct fan connection and full message control. "My thing was 'Why does this movie studio or record label own your entire fan base? You should own it and take it with you wherever you want and build whatever business you want on top of it,'" he says. Meanwhile, digital natives like fashion bloggers were moving to the app. "That's really where the influencer started to come from," he adds.

At the end of 2013, Beyoncé's team approached Porch about a surprise album release. "She doesn't want to do press; she doesn't want to tell anyone that it's coming. She just wants to come off the mountain and deliver an album," Porch recalls her team saying. "I don't even think Sony, her label, knew how she was going to launch it." She did so on Instagram, writing "Surprise!" to her more than 8 million followers as the full album appeared on iTunes. The successful rollout was proof that artists could sidestep the traditional press apparatus and other industry conventions if they had a strong enough connection with their fans. It was a flex of social media's power.

In the fall of 2014, the cover of *Vogue*'s September issue featured the nine hottest models of the moment, all of whom had gained large followings on social media. "Instagirls!" the cover line read. That was "the moment I really knew we'd tipped something in culture," says Porch.

"He came into tech at a time when they recognized they could become a supplemental engine to power in culture," says a tech adviser. "He, as a result, controlled massive marketing budgets to deploy that, and what that looked like was partnerships—funding parties, events." Porch came to serve as a kind of social fixer for Instagram co-founders Systrom and Mike Krieger. "I'm like, 'Great, you're gonna come to L.A.? I'm going to have Guy Oseary throw a cocktail party for you where you're gonna tell Ariana Grande about Instagram and why it's valuable,'" Porch recalls. "We're this new service. We're trying to convince people to use it. I want them to be able to picture the people that are leading it because otherwise they think it's going to be a bunch of robots that are controlling things behind the scenes."

When Adam Mosseri became head of Instagram a few years later, Porch ran the playbook back, getting his friend Jen Meyer, the jewelry designer who was then married to Tobey Maguire, to host a cocktail party for him at her house. Some of the Kardashians came. Porch became a trusted source for Mosseri on all things related to public figures and influencers and showed him the value of, say, Instagram sponsoring the Met Gala. "Charles is a generational talent at managing up," a former colleague of his at Facebook says. There are lots of people today who play this liaison role at tech companies, but Porch was a pioneer in the field, helping craft a template that Meta still follows. In 2015, he recruited Eva Chen, who, two years earlier, had been hand-picked by Wintour to lead *Lucky* magazine; she is now director of Instagram's fashion partnerships, one of several industry-focused subsections Porch created. "These are not socially gifted people, so if you have any social skills, you're ahead," says a former colleague of Porch's at Meta. "Technical engineers building products and Hollywood actors speak completely different languages. So for Charles to be able to explain them to each other in plain terms and not make either one feel stupid—he's very graceful in his diplomacy."

There was, of course, mutual benefit to these relationships. Porch's greatest skill, perhaps, is the fact that many around him tend to think otherwise. "You don't even know you're talking about work until it's over. It's almost like you leave the table and you're like, *Oh,*" says actress and influencer Dylan Mulvaney. "I think he's kind of a seed planter. The rest of his team are the ones executing in inbox and doing some of the heavier lifting, but he is so good at making you feel like you want to be a part of something." And, Mulvaney adds, "he's fun. That's the other thing. He's kind of the—I don't want to say 'life of the party.' I want to say 'a really good addition to the party.'"

When Porch moved to New York in 2019, he was approaching 40 and ready for a fresh start. Nine years in San Francisco and the city had never quite become home, in part because Porch traveled every weekend, most often to L.A., where he would crash at Foster's. He wanted more of a community and, after years of singledom, a partner. Also, Wintour told him to. "We were both at a wedding, and I get on the flight back—it was somewhere in the Caribbean—and my seatmate is Anna Wintour," Porch recalls. They knew each other only distantly through Instagram, but they talked the whole way home. "She's like, 'Charles, if you want to do culture, you need to move to New York.' And I was like, 'Well, if you say so.'" He bought a palatial apartment in an old townhouse in the West Village and began a yearslong renovation, documenting it every step of the way on Instagram, though he would end up living in the completed project for a mere six weeks. Because by then, he'd met Robert Denning.

Porch has seemingly always enjoyed proximity to power, not the public display of it. Denning, a philanthropist whose late father was one of the founding members of the private-equity firm General Atlantic, enjoys both. His Instagram is a catalogue of famous people he has met and events he has attended. But he is also an entertainer,

enrapturing dinner tables with his eclectic interests (aviation, for one) and hot takes. "Charles is such a professional, buttoned-up guy. He's very loyal, he's very thoughtful, but he's not, like, *Let loose and party.* He's almost got a work vibe, even socially," says a close friend of Porch's. "When we all met Robert, we thought, *This is amazing,* because Robert is not a filtered, edited guy. We were all just so happy for Charles that he could have that joy of life and the extravagance."

Denning first heard of Porch on a group chat with a bunch of gay men from Stanford business school. "I'm a reject of one of his friends," Porch tells me. One person in the chat wrote that he'd been told to go on a date with Porch—and he followed this announcement with multiple vomit emoji. Denning looked up Porch to find him chatting with Gayle King on *CBS This Morning* about onboarding the pope to Instagram. Says Porch, "He's like, 'I think he's cute. Can I go out with him?' The friend says 'no.' So then Robert's like, 'Now I *have* to go on this date.'"

Denning went on Instagram and found a mutual follower he knew from business school and Porch knew through Randi Zuckerberg. "And he DMs his friend and says, 'Becca, I'm sorry for not RSVP-ing to your wedding ten years ago or acknowledging your two children, but I need you to do something for *me*,'" Porch says. "And she writes back right away and she's like, 'I don't know him that well, but I'll try.'" The next day, Denning and Porch got drinks at Temple Bar. "Becca had a place of honor at our wedding, but we both barely know her," says Porch.

Denning had just moved to New York. He was recently sober and, following the fire sale of his luxury-eyewear start-up, trying to figure out what to do next. His father, Steve Denning, sat on several boards; as Steve's health declined, Robert started filling in for him, and his philanthropic career took off from there. In 2024, Denning became the then-youngest member of the board of the Metropolitan Museum of Art. "He's great on a board," Porch says.

By December 2022, they had made the relationship Instagram official, hard-launching with a photo in front of Joe Biden's White House (Denning is a big Democratic donor). And the following October, they were engaged. Denning thought they were going to a dinner at the Met in honor of an exhibit opening, but really they had the entire museum to themselves. Porch got down on one knee in the Greek and Roman wing.

In the spring of 2025, Amy Griffin, Foster and her two sisters, Naomi Gleit, and Meyer threw Porch a *Charlie's Angels*–themed bachelor party at the ultraluxury Amangiri resort, all of which was documented on Instagram. For the wedding, Porch and Denning hosted roughly 400 guests at Le Bristol, renting out the luxury hotel for the first time in its 100-year history. They paid for everyone's room as well as their room service. Night one was a party in the hotel courtyard; night two was the ceremony, held in the sculpture garden of the Rodin Museum. Ten minutes before the ceremony was to begin, the sky opened to a downpour, only to clear suddenly to an amber light. First came the bridal party, a group of six children—godchildren and spawn of close friends—outfitted by Thom Browne in seersucker suits, ties, and knee socks (Porch's idea). Then came the grooms, who walked down the aisle to the *Jurassic Park* theme (Denning's idea). Foster officiated. Barbara Bush gave the first reading, Griffin the last. The couple encouraged women to repurpose their old wedding dresses for the event; seemingly the only one who did was Real Housewife Carole Radziwill.

Then it was off to the Hôtel de la Marine for a reception and dinner. Before the guests sat down, Denning gave a vaguely political speech calling attention to the fact that this was a gay wedding and urging the roomful of people with lots of resources and power to take care of those whose rights they were celebrating. At the after-party, French drag queen Nicky Doll performed. The fairy-tale weekend ended with a black-tie affair at Château de Villette on the last Saturday of summer with the couple in custom midnight-blue Hermès tuxedos.

The pair have been well received in New York in part because they're so happy to be here. "Anna loved him because she could invite him to any dinner party and sit him next to any annoying or difficult person and he knew how to show up, lift the conversation, and never complain," one fashion editor recalls of Porch. "He's just very kind and normal and unassuming," says Wes Gordon, the creative director of Carolina Herrera and a friend of Porch's. "It's easy to forget he was a very senior person at Meta and now OpenAI."

Porch and Denning are tight with theater producers, fashion designers, art curators, and socialites such as Lauren Santo Domingo, Jenna Lyons, and Derek Blasberg. "They were two separately quiet people in different realms and then when they came together, they started making a splash," says one journalist. "A lot of people will make names for themselves in San Francisco and L.A. and then they come to New York and it's like being shot out of a cannon into other worlds."

A couple of months after the wedding, Porch was at Instagram HQ in Menlo Park when Altman's call came through. They discussed how "everyone's talking about AI and productivity and that's great, but no one's talking about AI and culture," Porch recalls. "And those two things are going to come together, and culture is run by all these creative industries, and question mark on what that's going to look like."

There were more conversations, with Altman and Fidji Simo, Altman's then-No. 2, with whom Porch worked for more than a decade at Meta. "I think what Fidji really realized is there was an enterprise-and-sales arm that was being built for it, but there wasn't anyone to work on the talent side," Porch tells me. "That is where there was an obvious opening for me to come in." Still, he had just gotten married, and he politely declined. When Simo came calling again in January, Porch was more receptive. "It was like, 'New Year, new me.' Time to jump to the next frontier," he says. "Sam, I feel, is the same person I met socially years ago in San Francisco. And it's kind of fun to work for a gay guy."

He started in March, spending his first few months on a listening tour, mostly in New York and L.A., meeting with everyone from the Council of Fashion Designers of America to HBO showrunners. "I ran around and went to a bunch of people I knew would have strong opinions on it, good or bad," says Porch. He went to the Cannes Film Festival, meeting with directors and hearing their feedback on the technology. "There was way more openness because directors have been using AI forever," says Porch. He also went to San Francisco—"a lot," he says, stressing the *t*—"to bond with the product folks and the researchers and try to get a sense of what they're really interested in. I don't have that many mathematicians on speed dial."

Then he brought people together for intimate events—"just a lot of relationship building for the company," says Porch. At Paris Fashion Week, he hosted a dinner at Sushi Park with Joopiter, Pharrell Williams's digital-first auction house, for various French creators. "It wasn't heavy programming, and it wasn't branded," Porch says. "It was just a collab with Joopiter to get people talking." He has also hosted dinners for *TBPN*, the daily tech-industry podcast OpenAI acquired in April to help bolster positive conversation around AI. The first *TBPN* dinner was held in L.A. and attended by the Foster sisters, Monica Lewinsky, *Crazy Rich Asians* director John Chu, and manager turned tech-investor Guy Oseary. They replicated it in New York a few months later at the downtown restaurant Kiko, where *TBPN* president Dylan Abruscato was seated next to *SNL*'s Chloe Fineman and Kareem Rahma of "Subway Takes." *Vogue* head of editorial content Chloe Malle and *Vanity Fair* editorial director Mark Guiducci were also in attendance. "He has been very instrumental in helping *TBPN* deepen our relationships with Condé Nast," says Abruscato.

A few days before our interview, I attended one such event, a dinner Porch co-hosted at the subterranean members' club Chez Margaux. In celebration of influencer Serena Kerrigan's new book, the gathering was attended by a mix of podcasters, creators, and Real Housewives, but, I noted to Porch, I didn't hear any discussion about AI. "That was a personal dinner, just so you understand. My personal and professional kind of blend," he says. "But what'll happen from that dinner is Serena's probably going to call me in two months and be like, 'Wait, I have this idea. I love Chat for this. Do you wanna do this thing with me?' So there's a nice cadence of how that happens. But also, you know, there were three Real Housewives there. If an AI model comes up on the show, it might be Chat because we have this nice drumbeat. This is what I've been trying to explain to folks: It's not about deals; it's about relationships."

Still, it's hard to say what exactly is supposed to result from this strategy. The idea appears to be that all this social lubrication will pay off and that Altman and his ilk will one day be seen less as freakish world destroyers than as parts of a broader mosaic of American celebrity. Familiarity, even at an elite level, tends to trickle down.

So far, though, things look decisively less glamorous than they did at Instagram. Aside from Lionel Messi's ChatGPT World Cup partnership, the most prominent celebrities working with the company have been rookie NFL quarterback Fernando Mendoza, who recently spoke to *GQ* about using ChatGPT for meal planning and nutrition, and designer Francesco Risso, who discussed the role of Astra in his creative process in an OpenAI-sponsored video on *Vogue*'s YouTube channel. "There might come a day when I get a call from Charles about doing something, and I'll navigate that in real time, but I recognize this brand has very bad PR right now," says one creative professional. They add, "If you're someone who doesn't like where things are headed, anyone affiliated with this world is going to be seen as part of the problem."

"When Kevin Systrom or Adam Mosseri were having events with celebrities, people wanted their leverage to help them grow their account, grow their influence," says one media executive. Instagram is inherently about marketing yourself; OpenAI's main product isn't yet public facing. "Aligning yourself with OpenAI, every star has to ask, 'What's in it for me?'—except for potentially a big check and a lot of blowback," says the media executive. Skepticism around technology has only grown since Hollywood cashed in on the NFT and crypto crazes, resulting in embarrassing commercials and promotions that blew up in many high-profile faces. And the threat AI poses to cultural fields also seems to be putting a higher premium on good old-fashioned human creativity and originality. "I don't care how cute you make Waymo look: This autonomous-AI-powered life is not warm and cuddly," says the media executive. "There's not a cool factor."

> "He's kind of the—I don't want to say 'life of the party.' I want to say 'a really good addition to the party.'"

At the same time, many in the entertainment industry are desperate. "Right now in Hollywood, only the top people are working," the executive says. "It's why Robert Pattinson is in every project or why Nicole Kidman is literally in everything you watch on TV. The next tier down is not getting jobs." According to Porch, when he joined OpenAI, there was a backlog of inquiries from various celebrities looking for an endorsement deal, production studios wanting to become more AI focused, fashion houses trying to figure out e-commerce, and sports teams and leagues looking to integrate more AI. It's possible that we will look back at this moment as we do at musicians who wouldn't give their songs to iTunes. "Now, it's like, 'I can't believe you would have turned down the check,'" one journalist says.

Some people roll their eyes at Porch, dismissing him as someone whose job is to be friends with famous people and questioning whether he has any expertise at all. "He's a super-wily leverager of social relationships, but it helps when you have multimillion-dollar marketing budgets to get people what they want," says a tech adviser. Another media executive sees what Porch has done as a kind of surfing: You choose your wave, start paddling at the right time, and rise at the perfect moment. It'd be fair to say Porch is a good surfer, they note. And now people are watching to see if he can stay up. "He sits in the middle of the convergence of media, culture, and tech, but he's also right at the changing of the tide," says the tech adviser, noting that Porch embodies a certain older-millennial brand that may be going out of fashion. "How much longer does someone like him maintain clout when it may be based on an outdated model of how those worlds interact?"

Porch insists he's not worried. "What I keep coming back to is, I think, the usefulness of AI," he says. "What it's transforming for people is undeniable. And the more conversations I have, I just—I feel it coming, people being able to talk about it more." He adds, "Everything moves faster in AI." ■

13.

NATHANIEL MORAN IS PIONEERING A REPUBLICAN-FRIENDLY POSITION ON AI SAFETY

→Congressman

→Age: 52

"THE ULTIMATE GOAL of what I should be doing in Congress is to increase liberty and the abundancy of life," says Nathaniel Moran, referencing John 10:10. "You can't have an abundant life, or you can't have a life full of liberty, if AI is running amok."

Moran represents Texas's First Congressional District, and his policies are moving what's acceptable in the GOP. At the end of June, less than one month before the Hugging Face incident disclosure, Moran introduced legislation to require frontier AI firms to report similar incidents. In July, he reached across the aisle to co-lead a piece of legislation that would require frontier AI firms retain the ability to curb or shut down their models, called the AI Kill Switch Act. "He was ahead of the curve, but the White House policy now is mandatory incident reporting," says Brendan Steinhauser, CEO of the advocacy group Alliance for Secure AI Action, which backed both of Moran's bills. "He's the leading Republican in the House on AI safety and security."

That's not to say he's breaking with his party. Moran backs AI laws that are narrowly tailored and pragmatic, calling himself a "solutionist." He also frames his AI policy in distinctly conservative terms: He often talks about AI as a matter of competition between global superpowers. He says he was introduced to AI policy serving on the House Judiciary Committee, when its IP subcommittee held a hearing on the use of AI in the creative arts. In preparation, he had ChatGPT brainstorm a list of questions and predict the witnesses' answers and was shocked to see that the chatbot was nearly spot on. "I would not have anticipated that it had that ability," he says. He is accordingly reluctant to pick a side between the increasingly bitter AI-safety lobbyists and their accelerationist counterparts, even if his policies are more favored by the safetyists. (When asked why he thought the pro-industry PAC Leading the Future was backing him, he said he didn't know but was "pleased" he had support from all sides.) Today, political operatives look to Moran as a signal of where his fellow Republicans, or even the president, may stand in just a few months' time. "We're going to be a leader in the AI space … whether or not that results in us having our name on a bill," he says. "And I can assure you no one will outwork us." V.I.

14-15.

Daniel Ziegler & Peter Lofgren Ushered in a Brand-new Donor Class

→AI-safety researcher; member of technical staff, Anthropic

→Ages: 31 and 37

STORIES OF POLITICALLY active tech workers are notable because Silicon Valley's workers are, by and large, apolitical. But in AI—partly due to many employees' more than $1 million pay packages, and partly due to the prevalence of effective altruism, which favors giving away much of that wealth—rank-and-file employees have begun engaging in election influence on an unprecedented scale. This election cycle, employees of Anthropic and OpenAI alone have given almost $5 million to political-action committees and campaigns. The bulk of that money backs candidates who have made AI safeguards cornerstones of their platform.

Peter Lofgren and Daniel Ziegler have become the poster children for the unprecedented political power of the average AI-company employee. Those who know Lofgren describe him as shy, a little socially awkward, and thoughtful—a quintessential engineer's engineer whose work helps make Claude refuse harmful queries. His largest donation this election cycle—just $100 shy of the $1 million contributed by Dario Amodei—was to Public First, an advocacy organization promoting strong guardrails on AI. He has also contributed the legal maximum to congressional candidates Scott Wiener, Alex Bores, Angie Craig, Jim Banks, and several others, totaling more than $1.4 million this cycle.

Ziegler has given slightly more than $1.1 million this cycle and has been spotted rubbing elbows with fellow donors and policy advocates at fundraisers. Those who know Ziegler describe him as endearingly nerdy and outgoing, and he makes himself available to lawmakers who want a read on what frontier technologists are thinking, sometimes through introductions from D.C.'s AI-safety lobbying groups and sometimes because politicians recognize him as a top donor. Most notably, Ziegler spent months as the lone donor to DREAM NYC, a super-PAC that backed Bores in his failed bid to represent New York's 12th Congressional District. After Ziegler's $50,000 seed donation, donors eventually contributed smaller amounts, and he was still the largest donor with contributions totaling $790,000. V.I.

16.

SNEHA REVANUR Is Writing AI Laws From the Outside

→Founder of Encode

→Age: 21

IN 2020, AT AGE 15, Sneha Revanur started the tech advocacy group Encode to fight a California ballot measure that would have replaced cash bail with potentially biased algorithmic risk assessments. Back then, she was worried about "AI being unreliable or error-prone in ways that would lead to super-consequential errors." After the advent of ChatGPT, that flipped. "My new concern," she says, "was that AI would be too powerful, too intelligent, too difficult for us to control." Three years later, three states—California, New York, and Illinois—have laws on the books requiring the biggest AI companies to publish and follow their safety plans and, in Illinois's case, submit to outside audits. "Every one of the state frontier AI laws that has been passed, Encode played a very major role," says Thomas Woodside, a fellow AI-safety-policy advocate. Those laws, in turn, are setting the baseline for federal AI regulation.

Revanur made Encode a political force while at Stanford. (She graduated this past summer.) In California, where advocacy groups can formally back bills, the group was an official sponsor of both SB 1047, Scott Wiener's ill-fated effort to regulate frontier AI systems, and its more successful successor, SB 53, last year. It also worked on New York's RAISE Act and Illinois's SB 315 with its staff successfully lobbying state lawmakers and governors to pass the legislation. On one occasion, a source tells us, Revanur flew out from California to meet Hillary Clinton in an effort to get her to back RAISE, which she did.

Encode has proved particularly adept at coalition building, something other AI safety organizations struggle with. "My impression is that folks in AI safety have been somewhat unwilling to compromise on their sort of epistemic snobbery about who they will and won't work with politically," says Revanur, whom colleagues describe as an optimist in an otherwise doomy scene. Encode, by contrast, manages to meet disparate groups where they are and look for common ground. It has rallied child-safety advocates, unions, and civil-rights groups together—most notably to kill Ted Cruz's attempt to override state AI laws last year. (That came just months after the group worked with Cruz on a deepfake-porn bill.) Now that she's graduated, Encode is her full-time job and has 15 employees. It's a nonprofit with funding from a range of philanthropies, including Prince Harry and Meghan Markle's Archewell Foundation. As Encode continues working on proposing legislation, Revanur has focused on trying to organize people inside frontier AI companies around AI-safety initiatives. She was one of the main organizers of this summer's "Pacing the Frontier" letter, a call from more than a thousand AI employees to consider slowing down AI development. S.H.



PHOTOGRAPHS: GETTY IMAGES (BORES, CLINTON); X (REVANUR)

17.

Alex Bores Is Already Shaping Democrats' 2028 AI Agenda

The defeated lawmaker has $30 million to keep presidential contenders focused on safety. By David Freedlander

→ Founder of
Who Decides

→ Age: **36**



IN EARLY NOVEMBER, in that exceedingly brief window between the midterm election and presidential-primary season, some of the biggest names in the Democratic Party are expected to gather at the Four Seasons in San Francisco for a closed-door conference on artificial intelligence. There, a number of presidential hopefuls will sit for a series of talks and dinners with leaders of the country's largest labor unions, donors to the biggest Democratic super-PACs, and political strategists who work with the party's top elected officials.

The host of the conference is not a former president or party leader or name-brand interest group; it is Alex Bores, a two-term state lawmaker from the Upper East Side who lost the Democratic

Artwork by Simon Denny

primary in June to succeed Jerry Nadler in Congress. Last year, Bores was mostly unknown in national politics. After sponsoring New York's RAISE Act in the State Assembly—which, as of January 1, will require AI companies to publish safety protocols and report incidents of potential catastrophic risk in their models—he campaigned on bringing similar regulation to the federal level. This prompted Leading the Future, a super-PAC that has received millions from an OpenAI executive, to spend roughly $8 million on an effort to defeat him.

The donation cast Bores as a martyr to the political power of tech companies. There were profiles in Politico, *Time,* and *Vanity Fair* and a 90-minute sit-down on Ezra Klein's podcast. Bores received a surge of campaign contributions as a rival set of AI companies and tech-world figures who support regulation poured $18 million into the race to counter LTF.

He ended up losing the primary to Micah Lasher, but in the process, Bores became perhaps the most preeminent voice for AI safety. It may seem like an odd fit: a politician with fame as an AI fighter who has major funders with ties to the industry. But the combination has helped Bores establish his next move. He is now launching an audacious campaign to push the field of 2028 Democratic presidential candidates to take on a new breed of corporate oligarch and write the rules for the most pivotal technology of the 21st century. "What's the David Foster Wallace essay about how a fish doesn't know what water is? We are going to be shaping the environment that these candidates are running in," Bores says. "We can't match Leading the Future tit-for-tat in spending, but that's all they have to offer. We will have public opinion and public support, and we will be shaping the terms of the debate."

Bores's *we* is Who Decides, an advocacy organization he started just months after losing his race, when he and his supporters were having conversations about how to maintain the momentum from his congressional campaign. (Who Decides raised $10 million at launch and has planned commitments of another $20 million.) Part of the work will be traditional advocacy. Bores has already met with many would-be 2028 contenders—among them Kamala Harris and Pete Buttigieg—or their advisers. But he also has a bigger idea in mind.

Who Decides is hiring dozens of staffers and placing them in the six early-primary states and the seven key swing states to meet with "grasstops" leaders—prominent local pastors, Democratic Party county chairs, heads of Indivisible groups. By organizing among powerful on-the-ground constituencies, Bores thinks Democratic candidates in 2028 will have to go on the record with their positions on AI policies to gain support from the party's base. "I want to make AI safety not like Medicare for All or 'Abolish ICE,' where half the people running are for it and half the people running are against it," he tells me when we meet in early October. "I want to make it like marriage equality or being pro-choice, where if you're a Democrat, of course you believe in AI safety."

Bores explains this while ordering a drink at an Irish bar on Murray Street. He'd just come in from a nine-plus-hour City Council hearing across the street that was headlined by Jacob Coxon, the former Anthropic employee who announced on social media that the technology his company was building could destroy civilization. After the meeting, a number of activists, operatives, and politicians had gathered to drink beer and talk about a "national Democratic AI agenda for 2028." The event, sponsored by Guardrails Action—another group working to push AI safety—was one of a series Bores has attended over the past few months.

As the Democratic presidential-primary race gets underway, party strategists increasingly view AI as the central tension of the contest, partly because AI touches on every other issue—wealth inequality, the economy, the environment, education. Concerns about data centers alone have morphed from a local land-use issue into a stand-in for the runaway power of tech companies and the possibility of an economy that makes most white-collar employment obsolete. And in part this is because it's impossible to talk about AI without acknowledging that Coxon's warning may prove prescient.

"Imagine the presidential campaign of 1944, and it's 1942, and instead of Los Alamos, you have Ford, GM, Standard Oil all fighting to create not a single bomb to drop on Japan but an entire nuclear arsenal. And imagine that in 1940, Charles Lindbergh had been elected. And imagine that you have Hitler and Stalin also racing the three biggest companies on earth to build their own nuclear bomb," says Daniel Squadron, co-founder of the States Project, which is dedicated to electing Democrats to state legislatures. "There is zero chance that ends well. And that's where we are."

So Bores finds himself tasked with getting Democrats to take on AI in a way that can appease both angry voters in the heartland and his big-tech donors. "We need to get to a place of 'yes, and' on these things," he tells me once his beer arrives. "It is not 'Oh, here's the hymnbook.' It is 'Let's all sit around a table and design the hymnbook.'"

When BORES RAN for Congress, he liked to point out that if he won, he would be one of the only members of the House with a computer-science degree. He has become known for talking about the possible harms of AI, but he often discusses the benefits, too. "I want AI that helps doctors diagnose disease instead of AI that allows health-insurance companies to deny claims," he told me when running in the primary. "This is a really potentially powerful tool, and we are on the cusp of having access to knowledge and power that we never had before."

As more voters have become aware of the threats of AI, some in the tech world worry that the views of the far-left flank of the Democratic Party will become its de facto position. Bernie Sanders has called for a moratorium on data centers and further AI research. Others on the left think AI companies should be treated as public utilities. "We have to use this moment where there is fury at this idea that a handful of billionaire tech bros—who aren't even offering us a product anyone really wants—might destroy the whole planet for their own trillions," Brad Lander, the progressive former city comptroller who ousted incumbent Dan Goldman in the June Democratic primary for Congress, tells the crowd at the Guardrails Action event. "There is fury, and that is useful."

Bores seems less keen to claim that no one wants the technology. His position is based more on oversight: requiring independent auditors, transparency on safety incidents, a ban on recursive self-improvement. Despite Bores's elevation as a go-to speaker about AI's worst outcomes, some Democrats think only someone who is explicitly anti-tech will suffice. "I don't care who you are: If you say something positive about data centers, it's going to be like saying on-camera that you wanted to abolish the police," says Morris Katz, the Democratic operative. "You have simply misread the moment, and you are never going to be president because of it."

David Shor, the head of data science at Blue Rose Research and something of a data guru to Democrats, has urged his party to notice that the national sentiment is not caution but rage. "People hate corporations. They hate the government, and they are extremely skeptical both of relying on the goodwill of these companies and of big government programs," he said at a recent conference, citing research that shows broad majorities—even among conservatives—favor heavily taxing AI companies to help people who have been negatively impacted by the technology. "Voters buy into this idea that if you lose your job because of AI, that is the government's problem and it is the government's role to find you a new job and to make sure that you can continue to pay your bills and not lose your home."

Bores's former backers are less bombastic in both tone and approach. "You can't go like Sanders," Chris Larsen, a crypto billionaire who, through a PAC, gave Bores $3.5 million for his congressional run and has emerged as one of the leading funders of the party, tells me. "That goes too far. It's just too simplistic, and it doesn't really have a plan. It's just sort of, *Oh, we can just put AI in a box and drop it in the ocean and it's all over.* You can't do that."

This tension has made many Democrats cautious. Some worry about looking foolish by being too anti-AI. The technology is changing rapidly; the politics could too. What happens if the AI stock bubble deflates? What happens if China races ahead? Then again, what happens if Democrats don't act fast enough? Liberals could lose the advantage on the issue after Donald Trump, who has catered to the AI industry, leaves office. Republicans like Steve Bannon and Missouri senator Josh Hawley have taken a much more aggressive approach to curbing the tech companies' power.

"What concerns me is that people really fucking hate this issue when you put it in front of them," says Rob Flaherty, a digital director for the Biden White House and deputy campaign manager for Kamala Harris in 2024, who has started his own super-PAC to support pro-AI-safety candidates. "And if we are weak and mealy-mouthed and then the Republicans do their demagogue thing—which would be about 'tech elites' and 'Big tech is coming to fuck up your community'—then that opens up a space for a conservative who is an economic populist."

Bores has promoted Who Decides as the solution: an outlet for voters' anger that is sensible enough to convince those who think like Larsen. But for his critics—especially supporters of Leading the Future—his attempt to find a middle ground only proves their point from the primary: Bores is not a legitimate voice for safety but a prop for one side of the debate. "He just got a $30 million bag from the same doomer groups who funded his campaign," Josh Vlasto, a co-leader of LTF, tells me of Who Decides.

"I want to make AI safety not like Medicare for All or 'Abolish ICE.' I want to make it like marriage equality or being pro-choice, where if you're a Democrat, of course you believe in AI safety."

A longtime New York political operative who came up in Chuck Schumer's office, Vlasto served as a top aide to former governor Andrew Cuomo. From there, he worked as a spokesman for Fairshake, the leading pro-crypto PAC. Vlasto maintains that whatever Bores may claim about taking on the power of the tech companies, his congressional campaign was supported by a different set of tech executives who warn of the perils of AI and call for strict controls because whipping up hysteria about the potential for doom is good for their share price.

"They are funding a narrative that is driving a threatening atmosphere to people who want to be thoughtful on these issues," he says. "And we want to stand up for people who want to pass a smart bill and smart policies in the face of relentless pressure."

Vlasto tells me this in his office in midtown as the dregs of a weekend nor'easter lash the windowpanes. When I arrived, he was barking into the phone, speaking with a client who runs a company that scans social-media to determine how much of any online conversation is coming from bots. A useful tool, he says, "when you are getting fucked with on social media."

When LTF launched, it was labeled a Republican operation masquerading as a Democratic one. And for good reason: The group's biggest funders are Andreessen Horowitz, the venture-capital firm that has given tens of millions of dollars to Republicans in the past few election cycles, and Greg Brockman, the president of OpenAI, who gave $25 million to President Trump's super-PAC last year. Vlasto's partner Zac Moffatt is a former Mitt Romney aide. LTF got involved in about as many Republican primaries this year as Democratic ones.

Despite this baggage, Vlasto says the notion that Americans oppose AI is a canard propagated by Bores's side and a mistake for Democrats. He cites polling data that shows less than one percent of voters mention AI and data centers as their No. 1 concern this midterm election cycle. While loud calls for banning data centers get the most play on social media, he says most Democratic politicians want data centers to go forward as long as meaningful guardrails are attached. So far this year, Vlasto says, almost no Democrats are running ads about banning data centers. If the Democrats go anti-AI, or anti–data center, he says, they risk alienating onetime parts of their base, such as blue-collar building-trade workers. "You go to Nevada, you go to South Carolina, and run a campaign that says 'AI is going to destroy the world,' and all these data centers that have created union jobs in working-class communities—I mean, I'm just not sure these guys have a feel for what the politics in this really are."

LTF'S COUNTERPART during primary season was the super-PAC network associated with Public First Action. The group, the biggest supporter of Bores's congressional campaign, was given $40 million by executives and engineers at Anthropic. Brad Carson, a former Democratic congressman for Oklahoma who serves as the super-PAC's leader, disagrees with Vlasto that his donors just want to write rules that favor their own employer. "Say whatever you want about Anthropic, but they're very sincere about wanting more regulation. They're not like making it up or pettifogging," he says. "I get why Josh is always saying this is an Anthropic-vs.-OpenAI fight. But if you're starting a super-PAC and someone calls you up and says, for whatever reason, 'We've decided we want to give $20 million to your cause,' you say, 'Okay, you have no impact on our governance. Sure, give us money.'"

LTF and Public First battled in congressional races elsewhere in the country in 2026, but the Bores race, held in the media capital of the world, got the most attention. Vlasto's playbook called for running a campaign that touched on all sorts of issues except for the one he was actually being paid to push. It backfired in Bores's race because of how many well-informed voters are in the district. "Manhattan is ground zero for global media, and when we started this campaign, AI wasn't even on people's radar," Carson says. "But all the media people live in the district, and suddenly Ezra Klein's covering it, Michelle Goldberg's covering it, and it makes AI issues front and center for the American elite."

Vlasto says Bores became LTF's first target because, after Trump's reelection, Anthropic was no longer Washington's preferred AI company, so the firm turned to the states to regulate. In Vlasto's view, Bores is no plucky upstart; he is a node in Silicon Valley's vast network of effective altruists. "I am not naïve. I think these fucking EA guys are going to put a billion dollars in a pot, and Alex Bores and like-minded groups are going to go up to each of the candidates and say, 'I have a billion dollars. This is the policy process we want you to follow,'" Vlasto says. "But I don't necessarily think it is going to work."

For LTF, beating Bores wasn't about just one race; it was about defeating a "doomer" worldview that Vlasto says wants to slow the pace of progress—and a rival company that would benefit from any kind of slowdown. Bores disagrees. "Their goal was not just to beat me but to crush me and to make it so that no one else wanted to run on AI," he says. "They failed horribly in both of their missions, and their donors have realized that."

In a company Slack message in June, Brockman, the OpenAI founder who gave LTF $25 million, part of a $50 million pledge to the group, announced he would withhold the second half of his donation, acquiescing to pressure from OpenAI engineers who disagreed with LTF's approach. Last month, a spokesman for Ron Conway, who was LTF's leading Democratic donor, announced he too would no longer fund the group. For Bores and Carson and the rest of their side, this was a cautious victory.

"I am sure these guys are going to reboot, and it's why you are seeing the pro-AI-safety wing get their shit in order," Jesse Ferguson, a longtime Democratic operative who advised Bores's congressional campaign, tells me. "These are the insanely super-wealthiest people in world history, and they are going to do everything they can to stop any regulation. And so we can defeat the super-PAC, but that doesn't mean we have defeated the money behind it—their first line of defense is going to be the Republicans, and their second is going to be the hearts and minds and bank accounts of the Democrats running for president in 2028."

"These are the insanely super-wealthiest people in world history, and they are going to do everything they can to stop any regulation."

Vlasto insists LTF is still moving forward. He says his group currently has $40 million in the bank. "This organization is funded as a multi-cycle organization, and I think we're just going to keep going," he says.

How this will play out in the 2028 primary remains uncertain. "It's a game of chicken," said one political operative who has been in communication with several would-be contenders. "No one wants to jump out ahead and be the AI-safety candidate. They know there is all this money up for grabs so they are all lying low and trying to be somewhere in the middle—not the best candidate for Silicon Valley but not the worst, either—and waiting to see where everyone else puts their stake in the ground. Leading the Future might be dead, but their money is very much alive."

For now, no Democrat serious about running for president is lining up as an outspoken defender of the AI industry. But Bores still says he is on the lookout for any indication that a candidate has gotten too close. "I will start knowing a lot more when their plans come out, but if there aren't hard limits in your plan, or if you are saying 'We'll keep talking about it,' or if you're doing dog whistles about how we have to beat China more than talk about how we have to protect Americans, and if you are giving this overly rosy story about how it will all work and only alluding to potential harms," he says, "that's a pretty good sign you are about to get laughed out of the primary."

And he is confident that his plan for how to shape a candidate's AI message will work. "They get their power from spending insane amounts of money and keeping advocates divided so they can pick people off bit by bit," Bores says. "But if all the voters are saying the same things in every town hall, and if every group that you call is saying it supports regulation, then the cost for buying candidates off is gonna go way, way up. And that is what we are making clear." ■

18.

Mike Forman Built a Data-Center Empire

→ Global head of digital infrastructure for
Blackstone Real Estate

→ Age: **36**

MIKE FORMAN STARTED looking at data centers in 2018 and 2019, long before ChatGPT sent the AI industry into overdrive. The Blackstone hypothesis was that businesses were moving their operations online at an extraordinary pace and demand for the buildings that made that possible would keep growing. He was one of three dealmakers (alongside Greg Blank, who leads the firm's data-center investments through its infrastructure business, and Tyler Henritze, head of strategic investments until 2023) who pushed for a $10 billion deal to buy data-center company QTS Realty Trust. When the deal was completed in 2021, Forman believed the market for computing infrastructure had much further to grow. But five years later, even he admits he underestimated how much the opportunity would explode: Blackstone reached its original seven-year investment projections for QTS in just seven months, and QTS is now 18 times larger than when Blackstone bought it, having expanded outside the U.S. It helped make the firm the world's largest data-center investor and Forman a trusted name within the company who can dictate where it splashes its cash—helping determine where the infrastructure that makes the AI revolution happen gets built and how quickly big tech can expand its AI arms. **C.S.W.**

19.

WHILE AMY KREMER Is Stoking a Grassroots Conservative Backlash Against Them

→ Chair of
Humans First

→ Age: **55**

AMY KREMER HELPED found the modern-day tea party, helped elect Donald Trump, and organized a rally on the morning of January 6. "No one is better at storming buildings and rallying people to extreme causes than Amy Kremer," Nathan Leamer, executive director of advocacy group Build American AI, said in May. He described her as "the Pierre-Augustin Hulin of America." In April, she was appointed chair of Humans First, a group incubated and spun out of nonprofit research firm Center for AI Safety that, under Kremer, has organized conservatives around AI risks.

Despite Trump's cozying up to the AI industry, the MAGA base of the GOP has resisted. Between March and July, the proportion of Republican voters saying they would oppose new data center construction in their area rose from 39 percent to 54 percent. Though 60 percent of data-center construction takes place in red states, most of the nationwide organizations stoking anti-data-center sentiment, like Democratic Socialists of America, are liberal. Kremer's is one of the only advocacy groups directly reaching out to the residents of deep-red towns and cities. Since her appointment, Humans First has organized a national day of protest against data centers and a bus tour this past September in which Kremer and partner activists visited 35 cities. One event in a 800-person Louisiana town, she says, drew more than 50 people: "People just kept coming and coming. The parking lot across the street started filling up with cars."

Her populist crusade has expanded beyond the centers themselves, which she describes as the industry's way of forcing artificial intelligence "down our throat" without public buy-in. She's invited parents of children who have self-harmed or died by suicide after bonding with chatbots to speak at events. At one in D.C., after a woman in the crowd identified herself as someone whose son killed himself after interacting with AI, Kremer began making arrangements. "Have you met with Senator Rand Paul?" Kremer asked the mother. "I will help you get Rand Paul, we'll help you get a meeting with Rand Paul."

Kremer is one of the few on the right who can tap into a pre-existing network of grassroots contacts to mobilize MAGA Republicans alienated by Trump's embrace of big tech, according to conservative lobbyists and organizers. "The people she knows from organizing the tea-party movement—there are people all over the country that she knows." **V.I.**

20.

THE COMPETITION BETWEEN THE WEST AND CHINA RESTS IN PART ON MISHA LASKIN

→ Co-founder of
Reflection AI

→ Age: **35**

THE RISE OF open-weight models has changed the economics of AI. Anyone can download, run, and modify them, rather than paying a company for access to its proprietary ones. By some estimates, more tokens now run through open-weight models than closed ones, but nearly all the popular ones come from Chinese companies, leaving a vacuum for a western-built competitor. The debate is part of a larger geopolitical battle over which country will own AI's future. Those in the industry who believe ceding the category to China isn't an option have found their best hope in Reflection, founded by two Google DeepMind researchers named Misha Laskin and Ioannis Antonoglou. Its investors include the Donald Trump Jr.–backed 1789 Capital; Nvidia also has a keen interest in Reflection's success, since more model-makers means more customers for its chips. It led Reflection's $2 billion funding round and remains one of its biggest backers. Laskin has, to many, become the face of the American open-weight movement. The undertaking is major; critics of Chinese open-weight models say they crib from OpenAI and Anthropic through a technique called distillation. Chinese labs can also lean on state-subsidized compute, which keeps their costs well below Reflection's. Reflection, meanwhile, has to build a model good enough to compete, give away, and still make money from enterprise customers. "I would almost say it's even harder to build Reflection than it is to build Anthropic," says one of the company's backers. **T. D.**

























Artwork by Simon Denny

21.

ALEXANDER BERGER WILL STEER THE UPCOMING FLOOD OF AI MONEY

→ Co-founder and CEO of
Coefficient Giving

→ Age: 36

MORE THAN A decade ago, the effective-altruism movement began preparing for AI to change the world. In 2015, Coefficient Giving, its biggest philanthropic funder, which has largely channeled the fortune of Facebook co-founder Dustin Moskovitz and has helped give away more than $7 billion, began nearly single-handedly spinning up the independent AI-safety ecosystem. It has funded many of the field's most prominent organizations, including MATS, Apollo Research, and Redwood Research. (It's also *Transformer*'s primary funder, though *Transformer* had no involvement in the writing or editing of this profile.)

But in recent years, while wide swaths of the public have come to share its concerns about AI, EA's reputation has cratered following a series of scandals, including, perhaps most notably, the fraud-driven collapse of the earn-to-give crypto exchange FTX. As a wave of AI employees unleashes newly minted fortunes, it falls to Coefficient Giving's CEO, Alexander Berger, to make sure his organization is still the best place for them to channel their philanthropy. He's done so by being the least weird among them.

Berger was slower than many of his colleagues to accept that AI should be the movement's central concern. He worried they'd become fixated on AI doom at the expense of other threats to humanity. In 2021, when he was appointed co-CEO of the organization, which was then known as Open Philanthropy, alongside long-standing CEO Holden Karnofsky, he focused the nonprofit's giving on other global crises: how to get bed nets to malaria-stricken regions, improve farm-animal welfare and global health, advance scientific research. This left Karnofsky to focus on the weirder portfolio: AI, biosecurity, effective-altruist community building. In 2023—a year before Karnofsky left for the Carnegie Endowment for International Peace, and two years before he took a job at Anthropic, which was co-founded by his wife, Daniela Amodei—Berger became sole CEO.

As Berger put it in a 2021 podcast interview, "I'm weirder than 99.9 percent of people, but not quite as weird as the last 0.1 percent." The economist Tyler Cowen sees the ascent of someone like Berger, his frequent interlocutor, as a sign of meritocracy. "If someone is leading a bunch of EAs, I want them to be pretty normal," he says. Some of the stranger elements of the ecosystem believe Berger's emphasis on normality is part of a strategy to distance Coefficient Giving from them. "Under that leadership, it became very clear the organization was much more focused on PR and reputation," says Oliver Habryka, chief steward of the rationalist forum LessWrong, whose organization Lightcone Infrastructure lost Coefficient Giving funding after he rented out his Berkeley event space for a conference featuring former white-nationalist writer Richard Hanania. "He does not view himself as a leader of the EA community, or the rationality community, or the AI-safety community."

Under Berger, Coefficient Giving has channeled more money than ever into AI safety. "As we started to see certain warning shots like the OpenAI–Hugging Face incident this year, he has consistently updated toward, 'We should be spending on the margin more resources to this,'" says Alec Stapp, a co-founder and the CEO of Institute for Progress, which received seed funding from Coefficient Giving. As Berger's focus has shifted to AI, the AI world has become more hospitable to people like him. In between panels at this year's Curve AI summit he stopped to chat with Neil Chilson, the head of AI policy at the Abundance Institute. "We did talk a little bit about how the weirdness is down a little bit." The way a friend of Chilson's put it in a group chat was it was less LessWrong and more LinkedIn. "The Curve this year was much more like Alex," Chilson says. "Maybe that's a good thing."

ANDREW FEDOROV

22.

Tom Brown

Is Tasked With Getting Anthropic on Trump's Good Side

→ Chief compute officer of
Anthropic

→ Age: 39

ANTHROPIC HAS spent much of this year at war with the Trump administration. First, it lost its contract with the DoD (and was earlier almost blacklisted across the government) after it refused to let its models be used for domestic mass surveillance and fully autonomous weapons. Then in June, its then-most-advanced models, Fable 5 and Mythos 5, were slapped with an export control order, effectively banning them. The order was eventually lifted, but briefings from within the administration against Anthropic and its leaders reportedly continued.

After the export control order, co-founder Tom Brown was sent in to make the company's case with the White House, convincing officials that Anthropic shared its aims of making its models safe. Formally, he's Anthropic's chief compute officer, responsible for getting enough chips and data centers to power its AI models. But as one of the few Republicans at the company, he also serves as a pseudo-diplomat: Brown joined CEO Dario Amodei at a White House AI summit in September. His soft skills and Republican credentials appear to make all the difference: "Tom Brown is not being a weirdo like Dario and can actually engage," one person told *Wired* earlier this year. There are signs he's doubling down on being the company's MAGA whisperer: *The Wall Street Journal* reported that Brown and his wife, Michelle Valentine, recently bought a house in D.C. and are positioning themselves as "allies of the Republican Party," according to its sources. The couple have donated to several Republicans this year, including $250,000 to the GOP's Senate Leadership Fund PAC. With his $15.5 billion net worth about to become a lot more liquid following Anthropic's IPO, Brown could well turn into one of the biggest Republican donors next cycle.

S.H.

WHO KNOWS WHO:

This year, Brown met with Elon Musk (whom he met at OpenAI, according to the 'WSJ') and agreed to rent compute from SpaceXAI, a $1.25 billion-per-month deal.



23.

BUT TINO CUÉLLAR IS THE COMPANY'S WHITE-SHOE ADULT IN THE ROOM

→ Chief global-affairs officer of
Anthropic

→ Age: 54

TOM BROWN STEPPED in amid a political crisis, but Mariano-Florentino Cuéllar is Anthropic's more deliberate bet on the near-term political future, one in which the current administration's power is less absolute and knowing how to call on Democratic networks will be important in shaping regulation. A former California Supreme Court justice with a professorial demeanor, he was appointed as Anthropic's first chief global-affairs officer amid the company's feud with the White House. He operates as the company's equivalent of secretary of State, managing government relations and representing a company more powerful than most nations to world leaders. He's tasked with guiding the policy proposals that the company puts before lawmakers, including the administration.

"Our general view is we need to develop new policy and get policy passed, and Tino has helped actually develop and advise on policy for many years," says a colleague at Anthropic. In a former life, Cuéllar served as president of Carnegie Endowment for International Peace, which gave him experience in "track-two diplomacy," the kind of relationships between institutions that get stuff done behind the scenes. There, he sat on Governor Gavin Newsom's commission on AI, which produced a report on AI safety that inspired California's SB 53, which at the time was the country's strictest AI-safety law. One former colleague said Cuéllar's greatest contribution was getting the other authors, which included a regulation skeptic and moderate alongside AI-safetyist Cuéllar, to cooperate. Many "were skeptical that a commission, or three co-leads, all with different backgrounds, were going to produce a report with coherent recommendations or policy principles," the former colleague explained. He's a "bridge builder."

V.I.

WHO KNOWS WHO:

While president of Carnegie, Cuéllar worked with Holden Karnofsky, the former CEO of Open Philanthropy and spouse of Daniela Amodei, who is the president of Anthropic.





PHOTOGRAPHS: GETTY IMAGES (DANIELA; MUSK); LINKEDIN (KARNOFSKY)

24.

INVESTORS ARE STILL BETTING ON LEOPOLD ASCHENBRENNER—EVEN AFTER HIS HEDGE FUND BLEW UP

→ Founder of
Situational Awareness
→ Age: 24

THERE'S A STORY Leopold Aschenbrenner likes to tell about his great-grandmother. She was a young girl when Allied forces firebombed Dresden, and, he says, the inferno was visible from her family's country home just outside the city. "It's valuable to have a sense of just how bad things can get," Aschenbrenner said to an interviewer in 2021. "I think this partially explains some of my proclivity, or interest in, existential risk."

Three months ago, things got very bad for Aschenbrenner. Over the course of a week, his hedge fund was forced to unload around $20 billion worth of publicly held stock, one of the largest trading losses by any investment firm in history. At 24, Aschenbrenner was already one of the most influential AI investors in the world. His two-year-old hedge fund, Situational Awareness, at its peak managed roughly $45 billion in assets. Day traders fetishized Aschenbrenner's insights to the point where, in some cases, Situational Awareness's public filings seemed to actually move markets. In May, the stock price for T1 Energy, a solar and battery manufacturer, popped after the firm disclosed a nearly $44 million investment in the company.

The $20 billion fire sale was shocking not just because of its size but also because Aschenbrenner, for a moment, was a familiar character: a boy genius who had convinced investors he was some kind of oracle. He'd cultivated this reputation in the three years since graduating from Columbia University, first at Sam Bankman-Fried's FTX Future Fund and then working on OpenAI's safety team. He'd gained notoriety with the paper from which his fund took its name, "Situational Awareness," which argued that AI's capabilities would soon outpace everyone's expectations by orders of magnitude that were hard for most people to comprehend. By then, he was surrounded by the most knowledgeable and influential people working in AI. Soon, he'd become a sort of John the Revelator, or as one podcaster called him, the "Nostradamus of AI." Someone built a website to track his various predictions. At a time when AI prophecy was in high demand, Aschenbrenner's paper offered a level of confidence that got investors' attention. "It was a seminal piece of content," says John Pfeffer, who runs a family investment office. Pfeffer reached out to Aschenbrenner after reading the paper and said he wanted to invest.

The natural next step was to start a hedge fund, and it wasn't difficult to find backers. John and Patrick Collison, the brothers who founded Stripe, signed on, as did Jane Street Capital and some of the leading AI researchers at Meta. The fund's early success—despite some losses, the firm ended 2025 up a reported 200 percent—was a credit to Aschenbrenner's audacity. At that moment, the global economy was making record bets on AI infrastructure build-out in the U.S., and Situational Awareness was taking enormous risk to maximize its bets, borrowing $3 for every $1 of its own capital. That meant the slightest wobble in the market could be ruinous. Which is what happened when Chinese companies released AI models that threatened the supremacy of U.S. companies' models. That, combined with sell-offs of AI and tech stocks, triggered margin calls from some of Situational Awareness' biggest lenders: Goldman Sachs, Bank of America, and JPMorgan Chase. In late July, while Aschenbrenner was preparing for his wedding—he was marrying Avital Balwit, his former co-worker at the FTX Future Fund and current chief of staff to Anthropic CEO Dario Amodei—he worked to come up with a plan to pay his lenders and save Situational Awareness from collapsing in on itself.

At first, he reportedly considered selling the firm's valuable stake in Anthropic. But instead, on the morning of July 30, he sold roughly $20 billion worth of publicly traded stock at a discount to Ken Griffin's Citadel. Investors haven't lost faith, basically arguing that Aschenbrenner's underlying bets on AI were sound, even if he'd overleveraged the firm. Aschenbrenner's rise was built on the conviction that he understood the near future of AI infrastructure better than almost anyone, and plenty of people are betting that thesis remains sound.

"Leopold learned a valuable lesson," Pfeffer tells me. "The risk was part of the bargain, and it wasn't existential because Leopold took decisive action." Situational Awareness still holds a stake in Anthropic, an investment that has likely already generated more gains for its investors than anything else in their portfolios. When the dust settled, after Aschenbrenner had unloaded $20 billion, Situational Awareness was still up over 80 percent on the year. Morgan Stanley, which was in talks to provide the firm with billions more in borrowing power before the meltdown, was reportedly not scared away.

A week after the fire sale, Situational Awareness quadrupled down on a long-term bet, investing $400 million in Source Foundry, a start-up Aschenbrenner had already backed that aims to make the machines and software that manufacture microchips used in AI systems. He recently hired several employees, including those who work on compliance (roles that might have prevented the company from overleveraging). But if Pfeffer worries about anything, it's that the $20 billion loss will chasten Aschenbrenner. "You don't make great returns without hubris, and you don't make outstanding, face-melting returns with box-ticking and bureaucracy. You do it through high conviction, edge, and aggression," he says. "The story of 2027 will be about the young genius's comeback."

JAMES D. WALSH

WHO KNOWS WHO:



Avital Balwit posted on X she'd asked planners to make the theme of her Carmel wedding to Aschenbrenner "European Garden, but it's on the cusp of the singularity." In attendance was Rob Granieri, the co-founder of Jane Street Capital, which had just lost $15 billion.

25.

And Ali Behrouz Might Teach AI How to Teach Itself When Humans Aren't Watching

→ Staff research scientist at
Google
→ Age: 29

RESEARCHERS HAVE yet to figure out how to get an AI to teach itself. That process—known as "continual learning" and popularized by Dwarkesh Patel—is seen by many as a crucial step to AGI: Right now, an AI model has to be retrained by developers to improve, but what if it could "learn on the job" as humans do, permanently accumulating new skills, task by task?

Ali Behrouz is "one of the most likely people to deliver" that breakthrough, according to Nathan Labenz, who interviews AI researchers for his Cognitive Revolution podcast. Behrouz's work on continual learning has produced some of the most important public developments to date, thanks to a few blockbuster papers that introduced new architectures and methods drawing on neuroscience to help models learn over time. The first paper, "Titans," gave AI systems an extra neural network, whose weights continuously update, to serve as a long-term-memory system. One follow-up, "Nested Learning," reframed AI models as layers of memory updating at different speeds and built a new AI architecture that learns how to update its own memory. A more recent paper, "Language Models Need Sleep," lets models consolidate their short-term memories by "dreaming," similar to how we do. As with so many other figures in AI, Behrouz views his work as more akin to ushering in a new species than developing a product. "It's a part of evolution," he says. "I can build something that takes over … or somehow be a representative of what I am. AI systems might be part of our society very soon" and "could have their own history."

S.H.

26.

ILYA SUTSKEVER

Could Reach Superintelligence First

→ Founder of
Safe Superintelligence Inc.
→ Age: 40

ILYA SUTSKEVER IS probably best known for his failed effort to oust his then-colleague Sam Altman from OpenAI, a saga memorialized in Luca Guadagnino's movie *Artificial*. A shame, since his contributions to AI are so much greater: He proved the power of neural networks back when most of the field had written them off (2012), helped build the AI system that beat world master Lee Sedol at Go (2016), and figured out how to use transformers to create powerful language models, laying the groundwork for ChatGPT and the AI boom that followed (2017). And now he's trying to achieve the field's ultimate goal: artificial superintelligence.

Sutskever co-founded OpenAI in 2015 and, for years, was its spiritual heart. "He has guru vibes," says one person who worked with him. Multiple others describe him as Yoda-esque, citing his tendency to speak in "oracular" aphorisms. "Success is guaranteed" was, according to Kevin Roose's *The AGI Chronicles*, a favorite in the early days of OpenAI; Sutskever led chants of "Feel the AGI" at the company's 2022 holiday party and burned an effigy of an "unaligned" AI the following year to visualize the importance of AI safety. But in November 2023, Sutskever, who was also an OpenAI board member, had become disillusioned with Altman and assembled a dossier of examples that he said exhibited a "consistent pattern of lying." Eventually, he and other board members tried to oust the CEO—failing in spectacularly public fashion.

In 2024, after leaving OpenAI, Sutskever founded Safe Superintelligence Inc., a company attempting to beat the big frontier companies from the outside lane in the race to develop superhuman AI, as in an AI model that outperforms humans at everything. To do so, SSI has eschewed the incremental releases of OpenAI and Anthropic and pledged not to release a single product before safe "superintelligence." Despite sounding like a terrible commercial idea, investors have lined up: He has raised around $8 billion at a $32 billion valuation, making it the best funded of all the "neolabs." Virtually no one knows what Sutskever is up to; the organization is incredibly secretive, discouraging employees from listing it on their LinkedIn profiles. And amid the constant churn in the industry, the team is apparently extremely tight, dedicated, and loyal to Sutskever. We're told that even Nvidia CEO Jensen Huang, who invested a reported $5 billion in the company, has not seen all its secrets.

S.H.

WHO KNOWS WHO:



Leopold Aschenbrenner dedicated Situational Awareness to Sutskever (who was his boss at OpenAI).



PHOTOGRAPHS: GETTY IMAGES (SUTSKEVER); X (ASCHENBRENNER, ALVITAL)