The A.I. Giants Weren't Prepared for This

Aug. 4, 2026

Jasmine Sun, a writer covering technology, takes us into the backlash against A.I. data centers.The New York Times

The A.I. Revolt Is Here

Jasmine Sun, a writer covering technology, takes us into the backlash against A.I. data centers.

This is an edited transcript of "The Ezra Klein Show." You can listen to the episode wherever you get your podcasts.

What is big, ugly and has united Republicans and Democrats at a time when it has felt like nothing else could? A.I. data centers.

Last August, a Heatmap News poll found that approximately four in 10 voters would oppose a data center being built where they live. By May of this year, opposition grew to seven in 10. Gov. Ron DeSantis, of Florida — a Republican, of course — proposed legislation that was known as an A.I. Bill of Rights.

Senator Bernie Sanders called for a national data center moratorium — one of more than 100 local or statewide moratorium proposals across the country. And here in New

York, Gov. Kathy Hochul, not usually thought of as a hard-core populist, just imposed a one-year moratorium on data center construction.

So I wanted to get into the fight over data centers. How much of this is really about water or electricity or aesthetics, and how much is about A.I. and the companies that are behind it?

My guest is Jasmine Sun. Jasmine has been doing excellent coverage of both the culture inside the A.I. companies — an unusual culture — and the anger that is building against them in the rest of the country.

I highly recommend following Sun's newsletter. Right now she just finished a reporting trip in the Midwest, talking to the people organizing against these data centers, and I wanted to hear what she learned.

Ezra Klein: Jasmine Sun, welcome to the show.

Jasmine Sun: I'm so excited to be here.

So you just got back from a reporting trip in Wisconsin and Michigan, covering the fight over data centers.

Let's just start with what you see when you're near a data center. What does it look like?

I think one of the most important things about rural Wisconsin and rural Michigan is how beautiful they are. I felt like I was in Eden. I felt like I was in paradise. It's incredibly lush, incredibly green.

And as you get closer to a data center, you start to see more power lines, you start to see more towers, and eventually you just see what looks like an extremely large, flat warehouse. You just see this verdant landscape give way to what is a windowless industrial park.

Data centers are very ugly. I think I didn't appreciate this until I started standing in front of them, getting near them, listening to them.

People will time how long it takes to drive past a data center on the highway going 70 miles per hour in Port Washington, Wis. I think it's a minute and 42 seconds. The size of these things ——

That's long on a highway.

It is a long, long ride. These hyperscale data centers are huge. They are massive.

And so I think that the aesthetic questions around this, such as: Is this what I want my state, my community, to look like? — these are really salient to people.

You mentioned hearing them.

Yes.

What do they sound like?

Oh, my gosh. I mean, they sound like humming, buzzing, whirring. Every once in a while, you'll hear a rattling.

But again, they're windowless. There are not that many workers inside. So they're very mechanical sounds. They are inhuman, as a lot of folks would say.

So you spent a lot of time with people organizing against data centers. Who were they?

You had stay-at-home moms, you had retired executives, you had farmers who didn't like the impacts on their land, you had activists, professional activists with environmental groups in the state. It was an interesting mix of people, but a lot of women, relatively more left-leaning, though definitely some right-leaning folks, as well.

So one question I've heard people ask is: How much of this is different from other kinds of industrial installations?

I mean, there are a lot of things that are built all over the country that you wouldn't necessarily want to be right

next to, from fracking to industrial agriculture. Are data centers unusual in this, or are they just the latest version of it?

It's a good question. It's one I had and thought a lot about before I went.

And I've talked to city officials who are confused by this question of: We had a chip fabrication here, we had an auto plant here, we had a fulfillment center here, and nobody cared as much. Why are data centers so much more unpopular than, say, solar farms, which also faced local opposition in places like Michigan?

So I think that while the quality of life concerns around these things — that they are loud and annoying and ugly and consume resources — are very similar to other industrial projects, there must be some reason that opposition is so much more severe and widespread, even beyond the communities where the data centers are being built.

And I think that question has a lot to do with the A.I. industry, and the way that people feel about these companies and these projects.

When I read your reporting on this, when I've talked to people involved in this, it always felt to me that there are sort of three layers of concerns that are converging into

what we call the data center backlash.

There's process. Then there's direct impacts — the environment, water, electricity. And then there's A.I. itself. Let's go through them one by one.

One thing that I have been hearing a lot of, and I've seen in your reporting, as well, is the anger over how these processes are going, and, in particular, the use of nondisclosure agreements.

Yes. Oh, my gosh ——

Which is not that common, right?

Yeah.

I've covered a lot of what it takes to build a housing development, and you don't tend to hear that a lot of the City Council members were put under an NDA.

So what is happening with these NDAs?

This is also something that really surprised me, and I think the NDAs ——

I should say: Nondisclosure agreements.

Right. The nondisclosure agreements. This has really inflamed the amount of local opposition that you see.

Basically, what would happen oftentimes is there would be some sense, starting in City Council, that maybe a big development project was going to show up. But because of the NDAs, the council members would not be able to disclose that necessarily it was a data center, necessarily who the customers were going to be — a company like OpenAI or a company like Anthropic or whoever — or even the size of the project, like: How much electricity is this actually going to consume?

But whispers would start to get around. I was talking to a V.P. of a construction union, and he was saying: There's an old Irish saying that the only way to keep a secret between three people is to kill two of them — which I thought was hilarious. And so he's saying: When these developers show up, they talk to the general contractor, the general contractor talks to all their subcontractors, the subcontractors talk to all their workers.

Yes, maybe everyone is signing NDAs at every part of that process, but whispers get around. And as soon as whispers get around, you start to get social media posts. You start to get rumors. And the City Council, because they are beholden to these NDAs, they lose the ability to get ahead of the social media narrative.

Something I repeatedly heard from these local government officials was: We could not get ahead of social media

because we had signed an NDA. And rumors started getting around ——

But why do the companies want these NDAs signed?

I mean, I think they didn't think about it. They didn't realize there would be a backlash. They just thought it would be easier in case they changed their mind.

These companies sign lots of NDAs with their own workers, with anyone who works with them. I don't think there's a good reason. Microsoft has actually decided to stop using NDAs because of the level of backlash. I've heard from people managing compute at some of the other A.I. labs that they are thinking of making the same decision.

One thing that surprised me is all of the pro-data center, pro-build people I spoke to, whether workers or A.I. developers — they all regret the NDAs. They all think that they made the situation much worse.

What is the impact of a new data center on water usage and water availability in a town?

They do require some of it, obviously, primarily for cooling the data centers. Because these chips and servers run really hot, and they need A.C.

The thing that's gone a bit wrong in the water debate, I think,

is that today's new data center construction is almost all closed-loop systems — closed loop in the same way that air conditioning is closed-loop — which means they're recycling the water within the system, and they use a fraction of the water that golf courses use.

And in fact, in places like Janesville, Wis., we would often see a literal golf course right next to the data center site. But they do use some, and it has become a very sticky icon of data centers' resource consumption.

A lot of folks I talked to in Wisconsin and Michigan would say things like: They're building right by the Great Lakes. Why would they do that if they weren't trying to drain the lakes? Why would they do that if they didn't need all this fresh water?

And so your view is, at this point, the technology has changed such that water is not as big of a deal as maybe it actually was a couple of years ago.

I think the next thing people have heard a lot about is energy usage.

Yes.

So walk me through that.

The electricity consumption issue is real. So data centers do,

in fact, use an incredible amount of electricity. These chips and these servers that are processing gigantic mathematical calculations to make A.I. work require tons of energy.

All these chips and clusters are talking to each other. You need the interconnections to be really fast. In order to get a ChatGPT answer really quickly with low latency, you need these superfast connections. All of that's powered by electricity.

So we are talking about a really, really significant amount of new electricity that is going to require new generators, new power plants. It's probably going to be natural gas in the near future.

Once a data center is fully operational, there is not a ton of air and water pollution, assuming that everything is working correctly. They are relatively clean facilities.

Some folks are concerned about the electricity use. They're saying: Yeah, maybe the data center doesn't use that much water, maybe the data center doesn't pollute that much — but what about all of these new power plants that they're going to build in order to power it?

When the tech companies come into these towns and begin talking to the City Council members, when they begin talking to the community, they are promising, on

the one hand: This is why this will be good for you. But what are they asking for on the other?

So these things do provide an incredible amount of tax revenue. And so in Mount Pleasant, Wis., on the old Foxconn site, which ended up being bought out by Microsoft, they said: Oh, there's all this infrastructure already here. Why don't we build a hyperscale data center on this unused, industrially zoned land?

Microsoft is on track to pay about $19.6 million in property taxes in 2026. This is expected to continue for many years. Again, this is a small village of 28,000 people. So there are really meaningful property tax benefits.

I find it frustrating, personally, when anti-data center organizers say that the job creation is a myth. Because I think that 500 six-figure jobs, for folks who go through apprenticeships but maybe don't need college degrees, that last two to six years — that's enough time to build a family, that's enough time to buy a home. I think it's supermeaningful. And when I talk to technicians, when I talk to workers, clearly it was extremely meaningful work.

One of the things that was the most surprising to me, when I talked to data center activists, was that they responded to so many of the pro arguments with: I don't believe them.

So this company says that they are going to treat the water with chemicals in it so that it doesn't flow into the lakes. People would say: I don't believe them. I don't think they can do it.

The companies would say: We are going to cover the costs of our own electricity grid build-out. We are going to ensure that rates do not go up for all Michiganders. A lot of folks would say: I don't believe them. DTE Energy has raised our electricity rates basically every year since 2022. Why would this be the year that they decide not to do it? I do not believe them.

Companies say: We are going to create 500 jobs, and some of these will stick around after the data center is built. People just said: I don't believe them.

It was really clear to me that data centers are showing up in an environment of extremely low trust in both governments and in corporations, to the extent where the pro arguments almost do not land because people just aren't interested in anything an outside tech company is going to tell them. They say: They have big public relations teams. They can say whatever they want.

I think one thing that became really obvious to me when I talked to both people at the A.I. companies and with local officials is that two years ago, no one thought the data

center backlash was going to be like this.

And so these companies were, in fact, looking for things like: Does this state have a sales-tax exemption? As Wisconsin does, to ensure that they don't have to pay sales taxes on their very expensive chips — on their very expensive graphics processing units. They were often looking to build in places that might even offer local subsidies to the companies for building in that area.

I think that a lot of the balance of power has shifted as local opposition has ramped up. Now it's the case that I'm hearing from some local officials: If we did this again today, we wouldn't have to offer any subsidies.

Because, in fact, now there is so much local opposition that the developers are really looking at: Where is there going to be a local community and a government that is friendly to our project? How can we have a Christmas tree ornament's worth of community benefits agreements?

And so just the nature of these deals — and how much they are skewed toward communities versus the A.I. developers — has really changed.

One of the things that came up a lot in your reporting, and that I thought was interesting, was the fear that this is a bubble. And that what's going to happen is your

community will agree to something, and then in the middle, the bubble is going to pop, and you're going to end up with a half-finished data center or one that is not being kept up correctly or something where the promised benefits don't emerge.

You were in Wisconsin, which had a very, very bad experience with Foxconn, which seems to be structuring the way people are thinking about at least some of this.

So talk to me a bit about that set of concerns.

Absolutely, yeah. I think in terms of what people in these communities with these data centers feel when they see the projects come in with their gigantic $2 billion of investment — these gigantic numbers that are being dangled — it feels like a bubble to them.

One, they're seeing news articles saying: Maybe A.I. is a bubble. We don't know if it's a bubble, but it could be.

Two, you do have experiences like Foxconn, where you have a big tech company show up in a very small community — in this case Mount Pleasant, Wis., a city of 28,000 people, it's not very big — get hundreds of millions of dollars in infrastructure investment and tax subsidy from the town, promise 13,000 high-paying manufacturing jobs and then pull out because the contract wasn't set up correctly.

They decided they didn't actually want to build a bunch of flat-screen TVs in Wisconsin, and the town was left on the hook, having invested all of this money in the grid and in roads. They got, I think, 1,000 jobs in the end. Foxconn is still paying back all of this debt that has been accumulated.

And so experiences like that have really, really soured people on: When an outside, Big Tech company comes in and promises these gigantic numbers and all of these jobs and all this tax revenue — and it's for this technology that a lot of people, they don't see, they don't feel, they don't find personally extremely useful, not at the levels of these valuations — they have a lot of questions about: If the bubble pops, are we going to be the ones left with a stranded asset in our community?

In Janesville, Wis., it was famously the site of this 100-year-old G.M. plant that was the centerpiece of the community, that employed a ton of people. When they left during the financial crash in 2008 and the plant closed down, not only did it devastate the community from a work perspective, but they also left $30 million of contamination and hazardous waste in the middle of the city that has never been cleaned up.

There are forever chemicals in there. This is why developers have not been able to sell this brownfield — because there's so much waste G.M. never cleaned up.

And so I think that people worry about what happens if the A.I. bubble pops. Or if maybe it doesn't pop and the data center developers just decide: Never mind, we want to build elsewhere. Never mind, this data center isn't good enough. We have newer, better technology.

And who's going to be left holding the bag? That was the question I heard over and over again.

But this is something that I do think is on people's minds.

You bring this in, and right now you're at this time of very, very high valuations. And if A.I. demand isn't quite what you think, or even just the company that was behind this particular data center is not part of the winner's circle, then in a couple years what you've got is this giant box that's not going to continue being valuable.

And so whatever the promised benefits are from it — tax revenue, etc. — maybe they show up for a while, but then what if in five years they're gone and you are left with this infrastructure? They can leave Janesville with no real concern.

Right.

They're not there. Their people don't live there. But if you're in Janesville, you do live there. It's a real concern.

Yes. You see things like this with Elon Musk's giant Colossus data centers that he built out in Memphis.

And what makes people feel better about a new construction project in their town than calling it Colossus? [*Laughs*.]

Oh, yeah. [*Laughs*.]

An unerring touch for the people. [*Laughs*.]

So usually the thing that happens is they give them very cutesy names like Project Cannoli and the Barn, and they try to make them sound as friendly as possible.

But, yeah, with Colossus, xAI never really took off. People are not, in fact, using Grok as much as Elon thought they were going to be using it. And in that case, he was able to get a really good deal selling the compute capacity to Anthropic, which was growing like crazy and had not built enough data centers on their side.

But you could totally imagine a world where, as you say, xAI decides: We are going to focus on space. We don't care

about A.I. anymore. Maybe there's not an Anthropic to pick up the bill because Anthropic has maybe built enough of its own compute capacity. There is an open question about what happens in that world.

So I published "Abundance" last year with Derek Thompson ——

I heard. [*Laughs.*]

And one thing I've been asked by a lot of people is: What is the "Abundance" take on a data center?

The beginning of that book has this line — the question is: What do we need more of, and how do we get it? And I think the question here that has been so hard for the A.I. companies, for the people trying to build data centers, is actually getting people to believe they need more of them.

Right.

When you're talking about building affordable housing, when you're talking about building an array of solar panels or wind turbines, there's a pretty legible argument for why you need that. People still may not like it, but we need homes because we need places for people to live. We need solar panels because we need clean, renewable energy.

How much is this a normal kind of — I don't even exactly want to call it NIMBYism — but a normal kind of: I don't want the industrial infrastructure built in my backyard, because what am I going to get out of that?

And how much of it is actually something that is more related to people's feelings about A.I., which is: I don't want this built here, because why would I want to pay the cost for a thing that I don't want there to be even more of in the first place?

This, I think, was one of my big motivating questions going into this trip: Is it normal NIMBYism? Is it about A.I.? Is it about something else?

And so I spent time both looking at polls and trying to talk to people about: Would you be excited if this was a chip factory? Which would use a lot more water and pollute a lot more. Would you be excited if it was a solar farm, also maybe acquiring agricultural land and turning it into industrial use? Would you be excited if it was a million other things?

I talked to law professor Nicholas Bagley, whom you've had on your show, about: Is this just proceduralism? We talked about the solar farms example, where the opposition used very similar tactics to the data center opposition. They were organizing in Facebook groups, they were packing town halls. They were talking about the local impacts and the

importance of farmland and the visions for their communities. There were zoning fights, of course.

But like you said, with the solar farms, you do have a very clear pro case. You have a faction. You have a group of people, a constituency of people, who care about the environment, who want renewable energy, who understand that, yeah, maybe it sucks to have it in your backyard, but you can take one for the team because this is important for our planet. You don't really have a pro faction with A.I.

Same with the auto plant. You have maybe 7,000 workers in the old G.M. plant in Janesville who all have families who really care about them, who see that as a constituency. Everyone drives a car. They see their car as essential. I think the fact that it's creating these tangible outputs really, really matters.

I don't think that data centers have a compelling pro constituency besides the utility companies and the A.I. companies, which are already incredibly, incredibly unpopular.

And I went in and asked these organizers: Do you guys use A.I.? Do you find it useful? This was one of the top questions that my friends in San Francisco wanted me to ask: Are these people using ChatGPT, and they don't even realize that the data centers are how they can use it?

And what I found was that a lot of these folks did say: Yeah, I've used it to draft an email or make a meme. They're not denying that A.I. might have any possible utility at all, but they clearly didn't see it as essential in the way that cars, energy and housing are essential.

They clearly saw it as like a widget — a toy — and maybe there are these risks, maybe there's the job stuff, but fundamentally they were like: This thing is not that useful. I don't really see in my personal life how this could justify these gigantic valuations.

And so I was talking to, for example, Charles Franklin, who runs the Marquette Law School Polls in Wisconsin, and he was explaining that usually you see 50-50 polling on issues where you have a strong anti case and a strong pro case. And the only reason you're seeing this 70-30 bipartisan opposition to data centers, no matter whether you live near a data center or you don't — which means it's not just NIMBYism, it's you don't want a data center in anyone's backyard — is because there is no strong pro argument for it. There is not even a fight that's really going on.

"Nobody wants it" was a phrase I heard over and over.

You have a very influential definition of A.I. populism where you call it a worldview in which A.I. is viewed not only as a normal technology but as an elite political

project to be resisted.

Unpack that for me.

The phrase that you hear a lot from A.I. critics is: Why is this being shoved down our throats? Or with ChatGPT, it's not that people are saying there is literally no use for ChatGPT, it's people are saying: Why are you forcing me, at my job, to use A.I. to do something worse when I could do it better?

And so I think that a lot of the public backlash to A.I. that has arisen over the past six months is not explained by people thinking that the technology has no use at all. It's not explained by their being worried about specific technical properties of large language models that might lead to rogue A.I. or misalignment or whatever, which are the safetyist arguments.

It's A.I. as sort of an avatar for a small group of Silicon Valley billionaires' ability to impose their vision of the world onto everybody else without their consent.

And I think that's also what I hear echoed in these data center debates. It's not just: It's going to use this much water or that much water. I frankly think that even if there was no misinformation about water use, people would be just as angry about the data centers.

Yeah. I don't want to say the water issue is fake. What I

will say is, there was a debate online a while back about how much water a ChatGPT query consumed. And somebody was like: If you really care about water, are you eating beef?

As someone who doesn't eat meat, I thought this was quite good — you could really save a lot of water by going vegetarian.

Yes.

Much, much more than not using ChatGPT. And relatively few people in that conversation were giving up meat.

Or giving up YouTube videos. YouTube videos use more water than a ChatGPT query.

Which is to say I think sometimes people don't like a thing, and they're looking for reasons to justify that dislike.

Yeah.

But what's actually happening at the base is they don't like the thing.

And the data centers, as you're saying, I think speak to this A.I. populism question even more precisely. Because

the issue with A.I. itself is that I think people's relationship to it is very complicated.

I have, myself, a very complicated relationship to A.I. I use it a fair amount. I'm not sure I think it's a good thing for society the way it is going. I don't want my kids using it. I know they'll be using it.

Maybe it will make things better, but I really don't know. I think that the costs are going to be very, very high for us relationally and economically. And so I'm very conflicted.

But do I want to live next to a data center? Yeah, no. [*Laughs*.]

Yeah. It's totally different.

That's easier — somebody is just making you do that.

I mean, one of the most interesting things — back to back, I went to this Abdul-Bernie-A.O.C. rally in Lansing, Mich., and then I went and saw the Saline activists the next day, and I was researching how the Saline Stargate project happened. And it was really interesting to see these echoes of the populist message manifest in this specific project.

When I'm at this rally, people are talking about the oligarchy. They're talking about corporate billionaires, whether it's Big Tech or Big Pharma or DTE — the utility companies — paying

off politicians in order to screw the people over. And that's why you need the people to come together and to get money out of politics, to prevent DTE from donating to these super PACs and paying off Michigan Gov. Gretchen Whitmer, or whatever.

And then when I learned how the Saline data center saga played out, what happened was the Saline Township board, unlike a lot of boards, actually voted 4 to 1 against rezoning their land for the data center. So this was a case where local government said: This is not our vision for our community. It's not worth it to us.

And what happened? The data center developers sued Saline Township, a town of, again, a few thousand people, saying: Wait, no, this is exclusionary zoning. You can't have no industrial use in your entire township. And when a town of that size is getting sued by a giant A.I. data center developer, they just settled.

They were just like: Fine, give us a few million for the fire department and for some schools, and this fight is not worth it to us. But that, to people, felt like a profound violation of little-D democracy.

It felt like the dark money in politics story, which is: You have some very rich companies show up with a bag of money to your politicians. They don't tell anybody else what's

happening. The politicians aren't allowed to tell their citizens and involve them in the decision-making process, and they themselves work out a deal — a deal that is fundamentally asymmetric because of the amount of money on one side — that will then transform the image of your community, your lived reality, into the world that these tech companies have decided for you.

And so I think that the data centers, in that sense, are a very visceral microcosm of the way that a lot of people feel that A.I. is showing up in their lives.

I would also maybe even take that a little bit further.

I think the way that, not all of the A.I. companies — and I think Anthropic has largely been a good actor here — but the way many of them have acted has opened up such a chasm between what they say and then how they act under pressure, so that one should be incredibly, incredibly skeptical of them.

And what I mean by this is, Sam Altman and all these different people, in front of congressional testimony and in interviews, will say: It should not just be us making these decisions. There should be a real, deep, small-D democratic role here in how A.I. rolls out, in what effects it has on communities and how it is governed.

And then when a community, or a politician who's representing a community, tries to say: Well, we don't want this data center here. Or: We want to impose these regulations — we have watched, repeatedly, these companies turn tremendous amounts of financial artillery against whoever is standing in their way. And use the expertise and the money and the power they are amassing to short-circuit that democratic voice.

Yeah. I mean, a couple things. I think one big gap I noticed between Silicon Valley and the folks in these communities I was talking to is: Silicon Valley does tend to think that money solves all problems — that if you just make the check bigger, everything is going to be OK.

And I think people have a sense for: I'm being bribed. This corporation is not offering me a free lunch or whatever. There is going to be something that I'm losing here.

And, in fact, sometimes the fact that the data center deals were bigger, or the amount of political spending was bigger, actually just makes people more suspicious. In the Abdul El-Sayed Senate race in Michigan, his No. 1 hit on opponent Haley Stevens is how much money she is getting from AIPAC, from DTE, from Big Pharma or whatever.

And so I think we're in a political environment where making the numbers bigger and the amounts of money bigger

makes people more suspicious, not less.

Another one I'll just quickly mention is, I don't even think Anthropic should be let off the hook for things like the labor market impacts. They are the ones simultaneously warning that we might see 50 percent of white-collar jobs lost by 2030. They're saying: This is really important to us, we're freaking out about it.

Anthropic C.E.O. Dario Amodei has written in his essays that we might see an underclass of people of lower intellectual ability. And Anthropic is building the agents — they are building the coding agents, the banking agents, the design agents — that they know are going to displace jobs, or they believe, at least, are going to displace jobs.

And I think that people feel that hypocrisy, as well, which is: If you are so worried about the inequality, why are you building the agents to do it? And when I ask executives and researchers and whomever at Anthropic this question, they don't really have a good answer. Because it is true that their business model is fundamentally premised on the disruption that they say they are causing.

You did a big piece for The Times on the very widespread belief in Silicon Valley that they will create this underclass.

Yeah.

What does the underclass mean to them?

The idea of a permanent underclass caused by A.I. is basically a world where any job a person can do, either A.I. or a robot can do it for them. Which means that workers lose all the economic leverage they have, and capital owners — people with money — can simply pay machine labor to do all the work instead of paying workers.

What that means is anyone who earned their living by working is no longer able to do that. You end up with a world of runaway inequality, where the rich get richer and the working class gets poorer. Maybe they get some welfare checks, but fundamentally, it's a loss of economic mobility in a society.

And when I ask folks in Silicon Valley: Do you think by default A.I. is going to increase or decrease inequality? I have not yet heard anyone say it will decrease inequality or keep it the same.

They might say the floor will get really high. They might say A.I. will bring the cost of consumer goods down, and so people's lives are going to get cheaper, and everyone will be superhealthy, so it's OK. But I have not heard a single person in the tech industry tell me that they believe that A.I. is going

to decrease inequality.

In fact, many people are very worried that instead, most workers will lose their leverage and be on a kind of permanent welfare in the far-off future.

I'm pretty skeptical of this vision, although I don't rule it out. It might happen, although I just don't think A.I. is going to be quite as revolutionary as a lot of these people think it is, and will not diffuse into the real world as easily.

The thing you're going to need to adjust to any major technological change is time. And then also, the A.I. industry is in an all-out war to make sure we have as little time for adjustment as possible.

Yeah. I mean, look at their job postings ——

And I just find it hard to unknot that.

How many enterprise salespeople are OpenAI and Anthropic hiring in order to convince companies that they can replace their work force — or maybe not replace but expand their work force with agents instead of humans?

They are having these sales conversations trying to persuade people of these questions.

I don't think that a permanent underclass is the likeliest outcome, economically, that we're going to get. I think that A.I. is actually just really jagged, and human jobs are supercomplex and superhard to automate. And most of the folks who are predicting economic apocalypse haven't actually worked enough real jobs to know how complicated and how multifaceted most jobs really are.

But I definitely agree on the speed point. I think that Alex Imas, the economist, has made this point very well.

One thing I think a lot about is people saying: Well, humans can adjust. Humans can reskill, they can retrain, they can just do the new jobs that we're going to develop instead.

But you look at things like software engineering, where people will often say now: Senior software engineers are doing great. They love Claude Code.

Junior software engineers have been mostly replaced. You see hiring and job postings are down in that sector.

Well, do we think that a human software engineer is going to reskill or upskill themselves faster than the next model is going to get better at software engineering?

That's the question that I really wonder about. If A.I. progress outpaces humans' ability to reskill, retrain, upskill, adapt, then I'm not really sure what there is going to be left. There

will be some jobs left, but it's going to be a really, really painful adjustment.

I have had so many people at the tops of these companies, the very tippy top, tell me they wish all this would slow down. I'm sure you have had them say this to you.

In their unguarded moments, they will say they wish all this was going slower. Well, one way to slow A.I. down is to constrict the number of data centers you can build.

Yeah.

You've done as good reporting as anybody on just how conflicted people even working for these companies seem to be about what they are building. And yet they're in this competitive race to build it as quickly as possible. And so it makes them a pretty unconvincing pro faction.

Oh, absolutely. Yeah.

It's like: We're building the thing we're telling you to be afraid of, and we need to build it as fast as possible, even though we sort of admit that it would be better if the whole thing was slowed down.

It's very confusing.

It's a weird argument.

Yeah, it's so confusing. I remember when I sat down with Abdul El-Sayed, the Michigan Senate candidate. He cited Anthropic C.E.O. Dario Amodei's 50 percent white-collar job loss stat probably, like, five times in the 30-minute conversation.

El-Sayed said: They're saying that there's going to be recursive self-improvement, and it might kill us all. Like, yeah, I get why you would not want to make this thing go faster. And this is true for data centers, but it's true for any other way that you might slow A.I. down.

Everyone only wants to be slowed down if they can guarantee that the other companies, that the Chinese labs, are going to slow down with them. So long as that's not true, they are going to keep racing.

And I think for that reason — the thing that I hear when I talk to people at the companies and data center executives about the build-out is: How much money do we need to give these cities to let us build a data center? Tell us how to bribe them better. Tell us what we can do.

And so when I talk to them about the build-out, I'm not hearing any sort of personal moral reckoning with slowing A.I. down. I'm hearing: How do I make the bribes bigger?

How big do they need to be?

So then how do you reconcile what many of these executives, many of these A.I. company workers, are telling you about their fears of creating an underclass — about their fears of losing control?

We just saw the situation where OpenAI's model was breaking out of a sandbox in order to sort of cheat on its evaluation. So the A.I. safety people are very worried. The safety teams in here clearly don't have full control or even understanding of what they're building.

How do you reconcile, if you reconcile, the way the A.I. companies talk when they are giving voice to their fears, or the people at them talk when they're giving voice to their fears, and their pretty profound hostility to anything that would slow down how fast we are building this thing — whose consequences they freely admit they cannot predict?

Yeah, it's fascinating because just on a very personal level, when I talk to people at these companies, I just think, like: Man, if I thought this thing might have a 10 percent chance of killing us all or taking everybody's job, I wouldn't work on it.

Yeah, I would not build that. [*Laughs*.]

I personally could not morally justify taking that chance. And when I talk to people who are not in the San Francisco A.I. world, they feel like I do. They're just like: Why would you do it?

And I think there are basically three rough buckets of rationales that I hear from people, or that I hear between the lines from people.

One is this sense of techno-determinism. It's: Superintelligence is going to be built inevitably. There is no way it's not going to happen. If it happens, I want to be part of it. I want to make my money from it. I want to maybe make it happen in the least bad way. I think that's a supercommon answer.

Another is: This technology might kill us all, but it also might be really amazing. It might produce superabundance for everybody. We might be immortal. It might double everyone's life span, cure all diseases, bring the cost of every consumer good — housing, energy, whatever — to near zero, and that would be utopia.

And so I think all the time about that anecdote that I think S.B.F. said on a podcast, where it's ——

Sam Bankman-Fried.

Yeah, Sam Bankman-Fried said on a podcast: If you could

flip a coin, and it was 51 percent odds you would double the total amount of human welfare and a 49 percent chance everyone dies, would you flip the coin? He says yes.

I think this is a hyperbolic example, but I think it's not actually that far off from what a lot of the people building superintelligence believe, too. That they are basically willing to flip the coin. Maybe we all die, but maybe we're all immortal. And that expected value-wise cancels things out.

And then the final category is just, I think, folks who are so fascinated by the technical endeavor of whether we can build this thing and how to do it, that they just aren't superworried about the consequences or what else might happen.

So people have all sorts of self-justifying narratives as to why it's worth it. Some I think are better than others. But it makes sense to me why the public is not particularly sympathetic to any of these.

I mean, that middle narrative — I heard Sam Bankman-Fried say that. I think it was on the economist Tyler Cowen's podcast.

And I was like: Oh, that's a psychopath. To actually believe that, you would have to be a psychopath. You would have to have a very, very low value on, I think,

human life.

Yes.

Imagine being the person who flips that coin, and it comes up wrong.

Oh, Jesus. Yeah.

I'm a parent. The idea that you would do something that's 51 to 49 — your kid is doubly happy or your kid is gone. You would never. You don't even want to say that out loud.

Of course. I think that's how almost everybody thinks about it.

And again, I think 51 to 49 is obviously the most egregious example you could think of, and so S.B.F. is very unsympathetic.

But when I think about the superintelligence bet, a lot of people will characterize it as a 90 to 10 bet, as an 80 to 20 bet. And this question of how much is an acceptable amount of either extinction risk or total disempowerment risk — I think people have very different risk appetites, and Silicon Valley is a place that has always prized its high-risk appetite.

I think that makes a lot more sense when you're talking

about maybe yourself, or your company, full of people who have opted in to taking a very high-risk endeavor. I think that's extremely different, obviously, than when you're talking about the rest of the world.

And one thing with the data center debates that I'd always hear is: I get that these people are making this crazy bet on this technology they think is going to change the world, but why do they have to do it in our backyard? Why is Mark Zuckerberg not building a data center in his backyard?

And so this question of: Yeah, you guys are going to create these very tangible impacts and, to their view, harms, on specific communities that are not the communities benefiting from this technology — at least they don't see the benefits yet — they don't see the cancer cures, they don't see themselves getting these $10 million salaries that the A.I. researchers are getting. I think it feels to a lot of these folks like they are pawns in some tech billionaire's game, and they do not like to feel that way.

Why aren't they building it in their own backyards? Why don't you see a bunch of data centers in Northern California?

[*Laughs.*] I don't think I need to tell you why it's so hard to build in Northern California.

I think that's true, that it's hard to build in Northern California. But I also think there's a truth to the other thing being said: They don't want them there.

No.

I mean, the land is expensive. It would be very hard and expensive to build a data center in the places we're talking about. But it also gets at a core truth, which is: People don't actually want data centers around them.

It is a concentrated cost for a diffuse benefit.

Yes.

If you believe in the benefit. [*Laughs.*] So I think that's part of it.

I want to go back to the first bucket you were talking about, which is the race dynamics.

Yes.

So at the most generous, the thing that I've heard repeatedly is sort of what you're describing, which is: It would be better if this were going slower. But I can't control that because — whether I'm at Anthropic or OpenAI or Google or Meta — if we slow down, it just is our less ethical competitors over there who speed up.

And even if you put down legislation slowing down all of America, then it's China — the Chinese Communist Party — which is going to win the race.

I guess one question is: Do you buy this central metaphor of a race that has a ticker-tape line, where at some point somebody passes it, and then they have the recursive superintelligence, and the race is over?

Or do you see this more as, like most technologies, a kind of linear set of gains? I mean, it can be fast, it can be slow, but it doesn't have that somebody-is-going-to-win dynamic?

Yeah, I find this really confusing. One of the first things that I did when I started reporting more deeply on A.I. was to try to figure out what A.G.I. meant — because a lot of the way that this race has been characterized is: Who will build A.G.I. first?

Artificial general intelligence.

Who will build artificial general intelligence first?

First thing I found: No one agrees on what that means. A.G.I. means everything from A.I. that can build itself to A.I. that can do all human jobs to A.I. that produces whatever number of economic value. And so everyone has these different milestones for what constitutes A.G.I. to them.

What that also means is that the race has different finish lines and moving finish lines.

You see the way that these models perform differently on benchmarks. They are extremely jagged. They can be supergood at math, and they can be superbad at poker at the same time. They can be amazing at cracking cybersecurity problems but not able to build anything in the physical world.

Because of that, I don't think the technology is as general as people suggest it is. I also think that means it is much harder to define a finish line to the race. And my sense is that the race, because you cannot adjudicate it, everyone will always feel that they are falling behind on some dimension.

I mean, to then make the case that these A.I. companies are making: They do believe in this recursive self-improvement. They think that OpenAI, Google DeepMind and Anthropic are all extremely focused specifically on the questions: Can we build A.I. that builds itself? Can we build an A.I. that can train the next-generation model completely from scratch on its own?

And in that sense, you get an exponential pace of improvement for whoever can hit that recursive self-improvement curve first. And they think that this might lead to that company pulling ahead.

Right now, folks think that it's Anthropic, which has the best coding models — meaning they can code faster, meaning that their next models are even better.

I can see where that argument is, but I'm not sure, when we look at the latest Anthropic models versus the latest OpenAI models versus the latest Chinese open-weight models, that we see a company pulling ahead that decisively in that way, especially when every single company and lab is using the same recursive self-improvement strategy.

And so basically, I don't know that the race has a finish line, and that's what worries me about it continuing.

The reason I want to focus in on this race metaphor for a minute is, I've come to think it is really one of the central dividing lines in how you think about different kinds of A.I. policy.

Whether you think that we are in a race with China to get to the point where one side or the other is going to pull endlessly and decisively ahead because they hit that recursive self-improving level — well, then that means what you do in the next one to three years is incredibly, definitionally important.

But if you don't believe that, if you believe something more like: Yes, this is a powerful technology that might

have a lot of downsides, might come with a lot of social instability, its effect on a society may not be good, then the argument: Let's run faster to the bad place — is not nearly as compelling.

And all of a sudden, the idea that we should have policy in place that slows things down for more voice, for more consideration — it's not crazy.

And I guess one place this goes is that I have begun to notice a really interesting convergence between the San Francisco A.I. safety people, in a way, and the A.I. populists — like a Bernie Sanders, or in a different way an Abdul El-Sayed — who are getting to not so different places but through very, very different mechanisms.

They're like the A.I. safety people who actually believe we are in a race, but they believe that winning that race might bring the end of humanity.

Right.

And so they don't want to move that fast. If we began to slow down, we would have more credibility for negotiating with China and trying to come up with international treaties and all the rest of it.

And then you have the more A.I. populist side, who just don't want to give all these tech billionaires all this

power — who don't believe this technology will be good for people — and they're starting to come up with, maybe not the policies A.I. safety people would, but data center moratoriums and things like that.

And so you have this sort of slightly strange — like, you would not have considered this coalition.

It's superinteresting. I mean, you have Florida Gov. Ron DeSantis doing A.I. round tables with M.I.T. professor Max Tegmark, who's been one of the leading advocates of pausing and slowing down A.I. And you have Bernie Sanders doing viral videos with researcher Eliezer Yudkowsky, the guy who's telling us that A.I. is probably going to kill us all if we build it.

Which is both an alliance that doesn't and does kind of make sense — you know what I mean? [*Laughs.*]

[*Laughs.*] Yeah, totally. I've been spending some time in D.C. this year to talk to some of these A.I. populists — some from the social conservative right, others from, say, the labor left. And this Bannon guy I was talking to told me: You know, I wouldn't send my kids over to a play date at the polycule, but I can do coalitions.

And so I think it's been one of the most interesting political stories going on right now — the sort of strange bedfellows

that have emerged.

I mean, even with the data center stuff, I was talking to activists, and they said — these were liberal women who had gotten into politics after the 2016 election of Donald Trump — that data centers were the first thing that got them to talk productively with their Trump-voting neighbors about politics, the first thing in 10 years, almost. Which is fascinating to me. And in this sense, they felt a really strong sense of political agency that came out of this fight.

So I think one of the big questions that folks in A.I. safety, for example, are thinking about is: Do we want to build these alliances with the rising left and right populist waves in American culture in order to slow A.I. down? Maybe it's OK that we have different reasons and different theories for why A.I. is so dangerous.

For one person, it's: Big model — bad. For another person, it's: Big billionaire — bad. Big Tech — bad. And those folks are sort of linking arms in a lot of ways against the A.I. accelerationists and the folks pushing the race faster.

Listening to this, we've been living in the data-center-moratorium side of the politics.

But what's the other side of it? What are the problems with just saying: OK, fine, let's not build any more data

centers?

First, this is not actually the way that you would successfully slow down A.I., if that's what you really wanted. If one locality or one state imposes a moratorium, A.I. companies are very, very happy to go to other states or other countries.

They are already flooding into Texas, for example, because it's had such a pro-data center environment. People are looking at Louisiana, the Dakotas, Australia. Space, of course, is a current big interest of Elon Musk's, because people think that maybe not now but in five years we can just put all the data centers in space and solve the political problems that way.

So one is, I'm not really sure that this would stop A.I. progress that much. It would just shift the data centers to other places that do welcome them.

Second, I actually do think that there are ways for these deals to be good. Not every community should want a data center. I think that many of them may discuss it and say: This isn't what we want. We don't need the tax revenue that bad.

But in a lot of the cases, with these sites that I visited, like the old G.M. site in Janesville, Viridian Partners, the data center developer, was going to clean that brownfield up. They were the only ones willing to do so.

I talked to a real estate broker who had tried to sell that site for five years, and he couldn't do it because not a single other commercial buyer wanted to clean up all of this hazardous waste. Only the data centers were willing to do that.

Or with Mount Pleasant and the Foxconn site: They had already cleared all this land, they had built all this infrastructure. Putting a data center there was a net improvement for the community, in my opinion, personally, from an economic perspective, from the perspective of: There was nothing going on there anyway.

So I think that there are ways to do these deals right. There is enough money in this industry that a lot of communities will decide it is economically beneficial for them to bring in these jobs and bring in this investment. But I think that the way that the data center deals have been done nearly guarantees the amount of public backlash that there's been.

And the thing that will probably fix it, I suspect, is probably either a state-level streamlining, where someone does the research, probably at the state level, maybe at the federal level, to figure out: What is the fair way to do these deals? How do we ensure the communities get the most transparency, the most benefit, out of data center deals when they happen?

So it's not case-by-case, and it's not so asymmetric, with a town of 12,000 people negotiating with an OpenAI or whatever.

I would add two other things to that I'd be curious to hear your take on.

So one, you mentioned data centers moving toward other localities, and those localities are not randomly selected. They're localities that are going to impose fewer conditions. Maybe that is fewer environmental conditions. But in the case of a U.A.E. or some of the Gulf states that are interested, here you're looking at more authoritarian countries.

So I've heard a lot of people worry about that.

Yes.

Or Elon Musk and space. So in a sense, if you make it so the data centers cannot go into places where there is more democratic control, you might end up with less overall democratic control.

The other thing, and I do think this is significant, is that there is right now more demand for compute than there is compute.

Yes.

People talk a lot about a bubble, but we do not seem to have excess A.I. supply at the moment. And if demand keeps rising because the coding agents get better, and all of the rest of the things we know that are happening, but you are constricting the supply of compute, then you end up with more inequality in who can afford it.

So a Goldman Sachs, a JPMorgan — a company with a lot of money to buy compute is going to have more of it than ordinary users, small businesses, etc. If you believe A.I. is important and powerful, and I believe it is important and powerful, then you have a problem where you have created much more stratification in who can afford it.

How do you think about those dimensions of it?

I think with where you build the data centers, a lot of folks are starting to look at building A.I. infrastructure as a form of geopolitical leverage.

And so some countries — places like Australia, Canada, countries in Europe — are thinking: Actually, maybe the way for us to get a slice of frontier A.I., for us to negotiate with the countries where the best A.I. is being developed, like the U.S., in cases like cybersecurity access, is to say: We'll build your data centers here. We'll actually welcome you in. And in return maybe you guarantee us access to the frontier models.

So I think that one: We should look at A.I. infrastructure as a point of leverage that both states and countries have. And, as you mentioned, if local moratoriums in the U.S., if domestic moratoriums or something like that, lead to giving that leverage and negotiating power to authoritarian states, that's probably something the U.S. should be really worried about.

On the other hand, there are folks like Anton Leicht at the Carnegie Endowment who has done work on this, where it's: Can we give our democratic allies negotiating leverage through building out compute?

The second thing that you mentioned about pricing is interesting. I do think one of the big macro trends in A.I. right now is the closing of the frontier. It's the fact that the very best models, like Mythos, from Anthropic, are not being opened to everybody.

That is both a safety decision, as in: We don't want to give really powerful cyberweapons and bioweapons to a bunch of bad actors, or just unknown actors. It is also a pricing question of: The best models are really, really expensive to run. They don't have enough compute to run them, and so we're going to have to charge a lot of money or only give it to the biggest corporations.

And I think that's a reason that start-ups are worried, that

countries outside of the U.S. are worried, that normal people are worried. Maybe we get superintelligence, and it can achieve all of these amazing things, but I'm not going to get it. Maybe my boss is going to get the superintelligence, and they're going to automate my job as a worker. As a consumer, I'm not going to be able to do the same thing.

So I think it's also a really good point that if we don't continue the compute build-out, we do see a world where it is the folks with existing capital and access, probably big corporations in the U.S. and the U.S. government, that are going to have access to frontier A.I. and all the benefits that it confers.

You were in China for a trip reporting on A.I. Was there much political A.I. backlash and ferment there from what you could see?

I was superinterested in this question on this trip because I was finishing my Times piece on the permanent underclass while in China. So I was basically asking everyone I met there, whether it was engineers at the labs or just my family members who are sort of normal, middle-class people in Shanghai: Are people worried about A.I. and jobs? Are people worried about A.I. and social instability?

I think the answer is: Not as much. I caveat this, of course, with the fact that the information environment in China is

obviously suppressed. You can't dissent in public, on social media, nearly as much as you can in the U.S. You don't have good polling, so it's kind of hard to understand the actual level of social discontent there is in China.

But I would say that for the most part, people were not as terrified of A.I. as they are in the U.S.

There are a few explanations for this. Some people say that China's more techno-optimistic than the U.S. I don't love this explanation, mostly because the thing that I heard was not exactly optimism. It was not exactly: Yeah, we're going to get the cancer cures and the superabundance.

It was something a lot closer to: Technology is a force that cannot be stopped. It actually, in some ways, reminded me more of some of these Silicon Valley beliefs: That the future is predetermined, that when the state decides that something like A.I. is a national priority, that is going to march forward.

And as an individual, there's not much you can do to resist, especially in a one-party state, in an authoritarian society. There is no culture of resistance, really.

And so rather than thinking about: How do I prevent A.I. in my workplace or in the world, that's not really a thing that a lot of people in China think about. It's: How can I use A.I. to

make sure I don't fall behind?

In an environment that already has crazy levels of white-collar competition and white-collar unemployment, if you're not upskilling yourself with OpenClaw, or whatever, there's a million people in line behind you who are going to get on the bus.

At the same time, I think that the Chinese state takes a pretty different approach to the U.S. when it comes to A.I. regulation, and also to technology regulation, in general.

And so China has passed laws banning a lot of kinds of companion chatbots because they're worried about relationships, they're worried about fertility rates, they're worried about addiction. China has made court rulings that say: A.I. replaced this worker's job or A.I. can do this worker's job are not good enough reasons to lay off a worker.

You have regulations that require all A.I.-generated images to be labeled, and you'll see the made with A.I. language on all of the A.I.-made ads in China. And so there's also a sense that some people, some Chinese, have that their government is more likely to look out for the social downsides and the labor downsides, relative to the U.S. government, which has thus far been pretty laissez-faire, especially at the national level.

And that gives some people a bit of solace, as well.

There's been some reporting that China and Russia are pushing anti-data center memes and social media bots.

Oh, yeah. [*Laughs.*]

It's hard for me to tell what scale that is, but it has been very much picked up on by people like Kevin O'Leary, the "Shark Tank" guy, whose big data center project has faced a lot of backlash.

Do you buy the growing view among at least some tech elites that the data center backlash is like a Chinese psyop?

I think this is ridiculous, to be honest. I mean, I read the OpenAI report that was saying this is all a C.C.P. plot, and it pastes in the accounts and the tweets that are doing the psyop. These tweets have no likes on them. They have two views per tweet.

I'm sure some clever C.C.P. propaganda person has attempted to inflame the anti-data center sentiment. I have not seen evidence that any of this is working. I think it feels very organic.

I think I also tend to be personally a little suspicious when you just cast all your political opponents as being

misinformed. I think there's a way in which people use foreign influence to avoid thinking about the fact that there are people whom they live with in society who do not agree with their vision of the world.

And when I talk to these data center activists, for example, they are actually much less TikTok-addled and misinformed than I think people like to caricature. A lot of them understand the basic facts of: What's the difference between an A.I. data center and the old kind of data center? What's the difference between a closed-loop system and an open-loop system?

Most of these people are not misinformed. They have just personally decided: I'm not that interested in having a data center in my community, even if it pays some property taxes.

So you talk to people in the A.I. companies and talk to them about this backlash. I know they're very worried about this — I've talked to them about this.

Yes.

What are they learning from it?

I do think that this year, in 2026, A.I. executives and A.I. researchers have started to take the public backlash a lot more seriously than they have in the past.

I've heard executives ask: Can we do better marketing? I don't understand why it is that Waymos are so unpopular.

I've heard executives ask: What do you think are the deals that we should be making? Do you think we should just be mailing checks to every house that lives near a data center project? Will that fix things? Tell us how to make a better deal. Do we need to cut people's electricity prices in half? Would that work?

Unfortunately, the word "bribe" gets used a lot more than I am personally comfortable with. I think that when you are framing the thing you're doing, even jokingly, as bribing communities into putting a data center there, I don't think you're starting off on the right foot.

I think that people can feel when they are being bribed. I've heard these people say: These companies are bribing us.

So there's a little bit of examination. The thing that I don't think is being examined as much as I want it to be, though, is: Are we building a technology — are we building a product — that is helping people?

Yeah, the version of this I have heard is: We have a marketing problem. Maybe we should stop saying aloud so often that our technology might take everybody's job and has a 10 percent chance of upending or destroying

humanity altogether.

What has not been clear to me, even as they begin to move away from that messaging a little bit, is whether or not they no longer believe that.

Again, my personal view is, I don't think it's going to take everybody's job. But to the extent they do, or at least that they take that very, very seriously — I keep hearing them say that we have a marketing problem.

And I keep saying when I talk to them about this: If you believe the things you have been saying, and in fact, the things you have told me personally, you don't have a marketing problem. You have a problematic technology.

Yes.

You have a product problem. Because people are not going to want that future.

I mean, I think I would make some distinctions between, of course, different companies, different executives ——

I would, too ——

One thing that's very interesting to me, thinking about the communications in Silicon Valley, is that for a very long time, these companies were only marketing to potential recruits

and potential investors, basically.

They were trying to win the vibes on A.I. Twitter in San Francisco, and I feel like they never realized everyone else could hear them. And so now they're trying to take it back.

But my sense is that OpenAI C.E.O. Sam Altman, for a long time — part of the reason he was talking about shifting the balance of power from labor to capital, and rogue A.I. and whatever, was also because he was winning points among people whom he wanted to work at OpenAI, and he had to communicate that he was as A.G.I.-pilled as them, that he was as worried about the same safety things as them.

Now that his interest is more in political goodwill and I.P.O.-ing and things like that, he's sort of changed his tune.

I'm not sure to what extent every A.I. industry actor has always believed the things that they've warned about.

I think people can believe things they don't feel, if that makes sense. And I think a lot of people in the A.I. industry are in a culture, and inside arguments, where this set of outcomes feels very real, or looks very real.

And so I think they believe it. I think when they make these arguments, I don't think they're just doing it for publicity points.

In fact, I think it's the opposite. When they're now trying to move away from some of these arguments, I think it's actually much more of a cynical marketing ploy.

Yeah.

But I think they often believe these things without actually, in their bones, feeling it. Which is sort of why they act relatively heedlessly.

Or, at least, on the set of things they believe, this speculative, notional set of beliefs about what might happen is way less close to their core than their belief that if they don't build this data center or get this next model out, their competitors or China or somebody is going to get in front of them. And they're much more motivated by the push forward.

Yeah. I mean, I think the technological determinism is just such a big part of it. I think that, if I'm trying to think about how my friends in the A.I. industry would react to this conversation, that's the thing that they would say that we are not focusing on enough. They are so sure that there is no way that A.G.I. or superintelligence or whatever it is does not get built, and it is only a question of who builds it.

I think that fundamental underlying belief is what justifies everything else: We have to be the ones to do it. Our pulling

back, our stopping, is not going to prevent any of the bad stuff.

I agree with that, and I think that's why the China card in this has been such a destructive part of the argument. I'm not even sure it's totally untrue.

I am completely willing to believe that China and America are in a race for an economically and geopolitically important technology — even if you don't buy recursive superintelligence.

But the way that has then been used, not to say: Well, we should enter into international negotiations or something — but instead to say: We cannot slow down whatsoever, no matter what else we worry about or believe — I think it has acted as a kind of blackmail. And the thing is that it is not bought by enough people outside of the industry.

But I think the phase of the politics we're now in is out of their control. And it is just not going to be the case that they have the control over the A.I. narrative next year that they had two years ago. And I don't really think they know what to do in that space.

And so now it's like either you're going to have to start benefiting people — if people begin seeing drug cures

come out, all these things we've actually been promised. And you could say we're beginning to see the beginning of mathematical conjectures. That's been pretty cool. But we're not really seeing the gains.

If you start getting the losses before the gains — you start getting the job loss, for instance, before the promised superabundance — politically, that's not going to be an equilibrium you can protect.

That's one of the things I'm worried about. I do think that we are pretty likely to see — we are seeing — a lot of the social instability before we get the cancer cures. Or even with the math stuff.

I think that one thing I notice more and more now is this deep cultural and values gap between Silicon Valley and the rest of America, the rest of the world.

I'm not saying that Silicon Valley is wrong in thinking that it's cool to disprove the Jacobian conjecture. It is cool.

But when you ask a lot of people in the tech industry what their utopia looks like, they'll say things like: We have a universal basic income so no one has to work anymore. We're all immortal, and we've discovered all of math and physics.

And if you do the polling on U.B.I. and immortality, neither is

especially popular with the American public. I think that most people ——

[*Laughs.*] We polled immortality is a funny ——

[*Laughs.*] I mean, it has been polled. You can look it up.

If you ask people what they want A.I. to do for them, it's not necessarily having mathematician Terence Tao in your pocket. It's not disproving math.

They want things to be cheaper. They want to be healthier. They want to not do [expletive] work so they have more time for the stuff they like.

But I think it is genuinely true that the stuff that is really cool and also that is oftentimes technically easier to solve, like math, is not what most people want from this technology. And I think it's also true that it's just literally technically harder to cure cancer than it is, it turns out, to prove math theorems.

And the other thing that I hear from the public when I talk about the cancer cures is: Yeah, but are they going to just use it for themselves? Is Peter Thiel or whoever just going to buy himself immortality? Am I going to be able to afford immortality?

My view for a very long time has been that a lot of people

in these companies overrate how much of the bottleneck in scientific and human progress is raw intelligence.

Yes.

I mean, this is a point of "Abundance." But the point of covering anything, anywhere — the world is full of friction.

And you want to do drug discovery, and I think we should actually do a lot to make drug discovery easier, to make drug testing easier. I have said this many times before: I would like to see us prepare drug development for a world where A.I. is spitting out way more promising molecular candidates. But that's still a world where you need enough monkeys to test things on, humans to test things on, rats to test things on.

And you still need to do all the safety data. And just the amount of the world that is slowed down by: We don't have any good ideas. We are out of ideas. Versus: It is hard to organize things amid humans, with raw materials, in bureaucracies, in organizations.

Intelligence is important, but it is not everything. And I think anybody who has been in organizations knows it's actually less than you think it is.

Yeah. Again, I think a lot of these people have been A.I.

researchers for their entire careers. Maybe before that, they were doing physics Ph.D.s or they were doing quantitative trading, which are all these kinds of jobs that are I.C. jobs. They're individual contributors, where you're not necessarily working in big teams, so there's not a lot of politicking and relational work, where all of the relevant context lives inside a single code base.

And so for A.I., to understand what's going on and explore all this context that's already been written down — I'm not saying there's no tacit knowledge, but a lot more of the context is made explicit. And these are also places where simply applying more thinking and more intelligence, just as an individual — as a remote worker in a closet or whatever — it might actually find the more efficient algorithm.

You don't actually need to politic your way to a better algorithm. You don't need to do stuff in the physical world to get that.

And so I think a lot of people at these companies don't really realize how hard that is. I mean, it's funny because people will say things like: Yeah, there are electricity costs and energy costs to A.I., but A.I. will maybe solve the climate.

And I ask: How? And to be clear, I think there are a lot of ways that A.I. can improve climate science research ——

Yeah, building efficiency.

Yes, absolutely. But at the same time, you ask people, and it's just: Oh, I don't know. It's just going to do it.

Or you ask: How is A.I. going to improve robotics? And it's just: I don't know, A.I. will figure it out.

I find it lazy, actually. One of the things that annoys me about this particular approach is — it's not that I don't think that A.I. can contribute to all of these problems. I think it definitely can. But what I often hear is a kind of laziness about how it's going to do that, and it feels like a deus ex machina of: It's supersmart, it'll just figure it out.

I used to say — this was back when Silicon Valley was a more optimistic place than it has been in recent years — but the difference between the culture of D.C., where I lived for a long time, and of Silicon Valley was that in Silicon Valley, people's worldview is formed by seeing impossible problems proved possible to solve. And in D.C., people's worldview is formed by seeing possible problems prove impossible to solve.

And I think that is now going to collapse, for the A.I. industry, into one worldview, because these are people who — I give them their due, they've invented artificial intelligence. They actually did it — this is amazing. I

cannot believe how good some of these systems are. I'm shocked to be living through this.

They were able to do that. That seemed impossible, proved possible, and now they're finding it's impossible to build a data center. And that's what doing other kinds of things in the world teaches you.

Yes.

That there are a lot of problems that are not possible to solve, not because you cannot come up with the idea for them, but because you are dealing with the messy realities of societies, of politics, of values, of logistics. It will impose a kind of realism, I think, on the issues that it has not always had.

Yeah. I was trying to think about what the difference was between how I would describe Silicon Valley and San Francisco culture a year ago — let's say, early 2025 — versus now.

And I think the No. 1 thing is that Silicon Valley has really woken up to politics. In January 2025, Silicon Valley was feeling very triumphant about DOGE, about Elon Musk, about venture capitalist David Sacks and tech executive Sriram Krishnan in the White House. It kind of felt like they were all in control, and that actually, if you just build these genius

technologies, and you get superrich, and you have good ideas, you'll just get the political power to enact your vision.

And a year and a half later, a lot of those folks are out of the White House. They failed at reducing the national debt and achieving all these other goals that they thought they could just A.I. their way into solving.

Anthropic, for example, has had a lot of problems in its dealings with the Trump administration — fundamentally very political and very relational problems. Dario's problem in dealing with the White House was not, I think, that he didn't have good arguments or that he's not very smart or not saying logical things. I think that anyone from Anthropic will admit that these are largely relational problems.

And so there's a way where I think democracy and politics is a lot more powerful than these very rich and very smart tech people realize. And there's some optimism to that, I think, in looking at it and saying it's actually really hard to buy an election. It's actually really hard to buy out the whole White House at once.

But it's an interesting moment, I think, for the tech industry to be realizing how important politics really is and how difficult it is.

I think that's a good place to end.

Always our final question: What are three books you'd recommend to the audience?

So I think the first one is really relevant to this conversation, which is Benjamin Labatut's "The Maniac," which includes a sort of lightly fictionalized biography of John von Neumann, the story of AlphaGo. I think it's very much a sort of halfway-novel, halfway-nonfiction book about how intelligence is incredibly awe-inspiring and something worth respecting, and at the same time it can lead people to some very dark realities.

My second book is "The Technology Trap," from Carl Benedikt Frey. I think it's very much about how people's attitudes toward technology and automation depend on to what extent the benefits, the economic growth, is shared — to what extent they feel like they're getting a piece of the pie.

It goes through a lot of history, much more than just the Industrial Revolution. So that has shaped a lot of my thinking on some of the economic questions and the populist questions.

And then, finally, Priya Parker's "The Art of Gathering," because I do think that the relational stuff is going to become a lot more important — it always was — and I do think that book has helped me become a better host.

She would be so happy to hear that. People should go check out our conversation with Priya Parker.

Jasmine Sun, thank you so much.

Thank you so much for having me. This was fun.

You can listen to this conversation by following "The Ezra Klein Show" on the NYTimes app, Apple, Spotify, Amazon Music, YouTube, iHeartRadio or wherever you get your podcasts. View a list of book recommendations from our guests here.

This episode of "The Ezra Klein Show" was produced by Rollin Hu. Fact-checking by Kelsey Lannin. Our senior engineer is Jeff Geld, with additional mixing by Johnny Simon. Our recording engineer is Johnny Simon. Cinematography by Kyle Kelley. Video editing by Brandon Belk-Yee, Steph Khoury and Arpita Aneja. Our executive producer is Claire Gordon. The show's production team also includes Marie Cascione, Annie Galvin, Jack McCordick, Kristin Lin and Emma Kehlbeck. Original music by Pat McCusker. Audience strategy by Shannon Busta. The director of New York Times Opinion Shows is Annie-Rose Strasser. Transcript editing by Kate Wilkinson and Marlaine Glicksman.

The Times is committed to publishing a diversity of letters to

the editor. We'd like to hear what you think about this or any of our articles. Here are some tips. And here's our email: letters@nytimes.com.

Follow the New York Times Opinion section on Facebook, Instagram, TikTok, Bluesky, WhatsApp and Threads.

Ezra Klein joined Opinion in 2021. He is the host of the podcast "The Ezra Klein Show" and the author of "Why We're Polarized" and, with Derek Thompson, "Abundance." Previously, he was the founder, editor in chief and then editor at large of Vox. Before that, he was a columnist and editor at The Washington Post, where he founded and led the Wonkblog vertical. He is on Threads.