

THE EZRA KLEIN SHOW

Why the Pentagon Wants to Destroy Anthropic

March 6, 2026



By Ezra Klein

Produced by Rollin Hu



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

Right now, as everyone is thinking about Iran, there is a story happening around the conflict that we need to not lose sight of. It’s about not only how we are fighting this war, but also about how we’re going to be fighting all wars going forward.

Last Friday, Defense Secretary Pete Hegseth announced he was breaking the government’s contract with the A.I. company Anthropic. And not only that — he intended to designate them a supply chain risk.

The supply chain risk designation is for technologies so dangerous that they cannot exist anywhere in the U.S. military supply chain. They cannot be used by any contractor or any subcontractor anywhere in that chain. It has been used before for technologies produced by foreign companies like China’s Huawei, when we fear espionage or losing access to critical capabilities during a conflict.

It has never been used against an American company. What is even wilder about this is that it's being used — or at least being threatened — against an American company that is even now providing services to the U.S. military as we speak. Anthropic's A.I. system Claude was used in the raid against Nicolás Maduro, and it is reportedly being used in the war with Iran.

But there were red lines that Anthropic would not allow the Department of War to cross. The one that led to the disintegration of their relationship was regarding the use of A.I. systems to surveil the American people using commercially available data.

So what is going on here? How does the government want to employ these A.I. systems — and what does it mean that they're trying to destroy one of America's leading A.I. companies for establishing conditions on how these new, powerful and uncertain technologies can be deployed?

My guest today is Dean Ball. He is a senior fellow at the Foundation for American Innovation and author of the newsletter Hyperdimensional. He was also a senior policy adviser on A.I. and emerging tech for the Trump White House, and the primary staff drafter of America's A.I. Action Plan. But he's been furious at what they're doing here.

Ezra Klein: Dean Ball, welcome to the show.

Dean Ball: Thanks so much for having me.

I want you to walk me through the timeline here. How did we get to the point where the Department of War is labeling Anthropic, one of America's leading A.I. companies, a supply chain risk?

The timeline really begins in the summer of 2024 during the Biden administration, when the Department of Defense, now the Department of War, and Anthropic came to an agreement for the use of Claude in classified settings.

Basically, language models are used in government agencies, including the Department of Defense, in unclassified settings for things like reviewing contracts and navigating procurement rules and mundane things like that. But there are also these classified uses, which include intelligence analysis and potentially assisting military operations in real time.

Anthropic was the company most enthusiastic about these national security uses, and they came to an agreement with the Biden administration to do this with a couple of usage restrictions: Domestic mass surveillance was a prohibited use in addition to use for fully autonomous lethal weapons.

In the summer of 2025, during the Trump administration — and full disclosure was in the Trump administration when this happened, though not at all involved in this deal — the administration made the decision to expand that contract and kept the same terms. So the Trump administration agreed to those restrictions, as well.

Then in the fall of 2025 — I suspect that this correlates with the Senate confirmation of Emil Michael as under secretary of war for research and engineering. He comes in, he looks at these things, I think, or perhaps is involved in looking at these things, and he comes to the conclusion that no, we cannot be bound by these usage restrictions.

The objection is not so much to the substance of the restrictions but to the idea of usage restrictions in general. So that conflict actually began several months ago.

As far as I understand, it begins before the raid in Venezuela on Nicolás Maduro and all that kind of stuff. But these military operations maybe increase the intensity, because Anthropic's models are used during that raid.

And then we get to the point where we are now, where the contract has kind of fallen apart and the D.O.W. — the Department of War — and Anthropic have come to the conclusion that they can't do business with one another.

And the punishment is the real question here, I think.

Do you want to explain what the punishment is?

Basically, the Department of War saying: We don't want usage restrictions of this kind as a principle. That seems fine to me. That seems perfectly reasonable for them to say no.

A private company shouldn't determine — Dario Amodei does not get to decide when autonomous lethal weapons are ready for prime time. That's a Department of War decision. That's a decision that political leaders will make. And I think that's right. I agree with the Trump administration on that front.

So I think the solution to this is, if you cannot agree to the terms of business, what typically happens is you cancel the contract, and you don't transact any more money. You don't have commercial relations.

But the punishment that Pete Hegseth, the secretary of defense, has said he is going to issue is to declare Anthropic a supply chain risk — which is typically reserved only for foreign adversaries.

What Hegseth has said is that he wants to prevent Department of War contractors — by the way, I'm going to refer to it variously as Department of Defense and Department of War —

I still call X Twitter. [Chuckles.]

Yes, I still call X Twitter. Right. It's just an inconsistency of mine.

Anyway, all military contractors can be prevented from doing any commercial relations, in Pete Hegseth's mind, with Anthropic.

I don't think they have that power. I don't think they have that statutory power. I think the maximum of what you could do is say that no Department of War contractor can use Claude in their fulfillment of a military contract. You can't say you can't have any commercial relations with them, I don't think.

But that is what Hegseth has claimed he's going to do — which would be existential for the company if he actually does it.

OK. There's a lot in here I want to expand on, but I want to start here. Most people use chatbots sometimes, if at all, and their experience with them is that they are pretty good at some things and not at others. And were not all that good in June 2024 when the Biden administration was making this deal.

So here you are telling me that we are integrating, in this case, Claude, throughout the national security infrastructure. It's involved somehow in the raid on Nicolás Maduro.

To what degree should the public trust that the federal government knows how to do this well with systems that even the people building them don't understand that well?

One thing is that you have to learn by doing. So it is the case that we don't know how to integrate advanced A.I. systems really into any organization, right? We don't know how to integrate them into complex pre-existing workflows. The way you do it is learning by doing.

Didn't Pete Hegseth have posters around the Department of War saying: "I want you to use A.I."?

[Laughs.] They are very enthusiastic about A.I. adoption.

Here's how I would think about what these systems can do in a national security context.

First of all, there's a longstanding issue that the intelligence community collects more data than it can possibly analyze. I remember seeing something from, I forget which intelligence agency, but one of them, that essentially said that it collects so much data every year that it would need eight million intelligence analysts to properly process all of it.

That's just one agency, and that's far more employees than the federal government has as a whole.

What can A.I. do? Well, you can automate a lot of that analysis — transcribing text and then analyzing that text, signals intelligence processing, things like that. That's one area. Sometimes that needs to be done in real time for an ongoing military operation, so that might be a good example.

Then, another area is that these models have gotten quite good at software engineering. So there are cyberdefense and cyberoffense operations where they can deliver tremendous utility.

Let's talk about mass surveillance here, because my understanding from talking to people on both sides of this — and it has now been fairly widely reported — is that this contract fell apart over mass surveillance at the final, critical moment.

Emil Michael goes to Dario Amodei and says: We will agree to this contract, but you need to delete the clause that is prohibiting us from using Claude to analyze bulk-collected commercial data.

Yes.

Why don't you explain what's going on there?

The first thing I want to say is that national security law is filled with gotchas.

It's filled with legal terms of art, terms that we use colloquially quite a bit, where the actual statutory definition of that term is quite different from what you would infer from the colloquial use of the term. Things like “private,” “confidential,” “surveillance” — these sorts of terms don't necessarily have the meaning that they do in natural language.

That's true in all laws. All laws have to define terms in certain ways that are not necessarily how we use them in our normal language. But I think the difference between vernacular and statute here is about as stark as you can get.

“Surveillance” is the collection or acquisition of private information, but that doesn't include commercially available information. So if you buy something, if you buy a data set of some kind and then you analyze it, that's not necessarily surveillance under the law.

So if they hack my computer or my phone to see what I'm doing on the internet, that's surveillance.

That would be surveillance. If they put cameras everywhere, that would be surveillance.

But if there are cameras everywhere, and they buy the data from the cameras, and then they analyze that data, that might not necessarily be surveillance.

Or if they buy information about everything I'm doing online, which is very available to advertisers, and then use it to create a picture of me — that's not necessarily surveillance.

Or where you physically are in the world. Yes.

I'll step back for a second and just say that there's a lot of data out there, there's a lot of information that the world gives off — your Google search results, your smartphone location data, all these things.

The reason that no one really analyzes it in the government is not so much that they can't acquire it and do so. It's because they don't have the personnel. They don't have millions and millions of people to figure out what the average person is up to.

The problem with A.I. is that A.I. gives them that infinitely scalable work force. Thus, every law can be enforced to the letter with perfect surveillance over everything. And that's a scary future.

We think of the space between us and certain forms of tyranny, or the feared panopticon, as a space inhabited by legal protection. But one thing that seems to be at the core of a lot of fear is that it's, in fact, not just legal protection. It's actually the government's inability to have the absorption of that level of information about the public and then do anything with it.

Yes.

And if all of a sudden you radically change the government's ability without changing any laws, you have changed what is possible within those laws.

You were saying a minute ago that “mass surveillance,” or “surveillance” at all, is a term of legal art, but for human beings it is a condition that you either are operating under or not.

The fear, as I understand it, is that either the A.I. systems we have right now, or the ones that are coming down the pike quite soon, would make it possible to use bulk commercial data to create a picture of the population and what it is doing.

Then the ability to find people and understand them goes so far beyond where we've been, that it raises privacy questions that the law just did not have to consider until now — so the laws are not up to the task of the spirit in which they were passed.

I would step back even further and say that the entire technocratic nation-state that we currently have in the advanced capitalist democracies is a technologically contingent institutional complex.

The problem that A.I. presents is that it changes the technological contingencies quite profoundly. What that suggests is that the entire institutional complex is going to break in ways that we cannot quite predict. This is a good example.

Not only is this a major and profound problem, but it is an example of a major and profound problem of a broader problem space that I think we will be occupying for the coming decades.

What do you mean by technological contingencies?

Well, the current nation state could not possibly exist in a world without the printing press, in a world without the ability to write down text and arbitrarily reproduce it at very low cost. It couldn't exist without the current telecommunications infrastructure.

The nation state is built dependent upon the macro-inventions of the era in which it was assembled. That's always true for all institutions. All institutions are technologically contingent.

We are having a profoundly technologically contingent conversation right now. A.I. changes all of this in ways that are hard to describe and kind of abstract.

This thing that we call A.I. policy today is way too focused on what object level regulations we will apply to the A.I. systems and the companies that build them — instead of thinking about this broader question of: Wow, there are all these assumptions we made that are now broken — and what are we going to do about them?

Give me examples of those two ways of thinking.

What is an object level regulation or assumption? And then what are the kinds of laws and regulations you're talking about?

An object level regulation would be to say that we are going to require A.I. companies to do algorithmic impact assessments to assess whether their models have bias. That's a policy I've criticized quite a bit, by the way.

You could say: We're going to require you to do testing for catastrophic risks. Things like that. That's an important area that we need to think about, but that's just one small part of the broader issue that our entire legal system is predicated on: imperfect enforcement of the law.

We have a huge number of statutes, unbelievably broad sets of laws in many cases, and the reason it all works is that the government does not enforce those laws anything like uniformly. The problem with A.I. is that it enables uniform enforcement of the law.

Here's the Pentagon's position: They're angry at having this unelected C.E.O., whom they have begun describing as a woke radical, telling them that their laws aren't good enough and that they cannot be trusted to interpret them in a manner consistent with the public good.

Secretary of Defense Pete Hegseth tweeted, speaking here of Anthropic: “Their true objective is unmistakable. To seize veto power over the operational decisions of the United States military. That is unacceptable.”

Is he right?

I have not seen any evidence that Anthropic is actually trying to seize control at an operational level. There’s an anecdote that has been reported that apparently Emil Michael and Dario Amodei had a conversation in which Michael said: If there are hypersonic missiles coming to the U.S., would you object to us using autonomous defense systems to destroy those hypersonic missiles?

And apparently Dario said: You’d have to call us.

I have been told by people in that room that is not true.

I have been told by people in that room that did not happen.

Not only that, but that there was a broad-speaking exemption for automated missile defense that would make that irrelevant.

That’s exactly right.

I am worried that there was a lot of lying happening here by the Trump administration.

Look, I think that’s probably true. I think that there’s lying happening, too, to be quite candid.

I don’t think that Anthropic is trying to assert operational control over military decisions. That being said, at a principal level, I do understand that saying autonomous, lethal weapons are prohibited feels like a public policy more than it feels like a contract term.

So it does feel weird for Anthropic to be setting something that does, if we’re being honest, feel like public policy. I don’t think it’s as beyond the pale or abnormal as the administration is claiming. And one way you know that is that the administration agreed to those same terms.

I think this gets to something important in the cultures of these two sides.

Anthropic is a company that, on the one hand, has a very strong view — you can believe their view is right or wrong — about where this technology is going and how powerful it is going to be.

And compared to how most people think about A.I. — and I believe that is true even for most people in the Trump administration, who I think have somewhat more of an A.I. is a normal expansion of capabilities view — the Anthropic view is different.

The Anthropic view is that they're building something truly powerful and different. They also have a view of what their technology cannot reliably do yet.

Some of their concerns are simply that their systems cannot yet be trusted to do things like lethal, fully autonomous weapons — which I don't think they believe in the long run should never be done.

Yes.

But they don't believe it should be done given the technology right now, and they don't want to be responsible for something going wrong.

On the other hand, they believe that they're building something the current laws do not fit. The view that Dario, or anybody, wants to control the government — I don't think Dario should control the government.

But I'm very sympathetic to — if I built something that was powerful and dangerous and uncertain, and the government was excitedly buying it for uses that could be very profound in how they affect people's lives — wanting to be very careful that I didn't sell them something that went horribly, [expletive] wrong. And then I am blamed for it by the public and by the government. That just seems like an underrated explanation for some of what is going on here, to me.

I think this characterization is accurate. I come out of the world of classical liberal think tanks, the right-of-center libertarian think tank world. That's my background. So deep skepticism of state power is in my DNA.

It's always funny how it turns out when you just apply these principles. You will sometimes end up very much on the right, and you will sometimes end up on the left. Because these principles transcend any sort of tribal politics.

This is like: No, we actually need to be concerned about this, and I think it's not crazy.

If I were in Dario's shoes, personally, I don't know that I would have done the same thing. I think what I would have done is said that contractual protections probably don't do anything for me here.

If I'm being a realist, if I give them the tech, they're probably going to use it for whatever they want. So maybe I don't sell them the tech until the legal protections are there, and I say that out loud. I say: Congress needs to pass a law about this. That would have been the way I think I would have dealt with it.

But again, it's easy to say that in retrospect. And you have to acknowledge the reality that what that means is that the U.S. military takes a national security hit. The U.S. military has worse national security capabilities.

Or they work with a company you trust less. Given that Anthropic has always framed itself —

But no company wanted this business. No other company did this.

Somebody was going to want it soon.

Someone was going to want it eventually, but no one took it for two years.

I think Elon Musk would have happily taken it over the last year.

Sure.

I've been curious about why Anthropic rushed into the space as early as they did.

They didn't need to do that. That's sort of my point.

In general, one of the odd things about them is they're people who are very worried about what will happen if superintelligence is built. And they're the ones racing to build it fastest.

A general, interesting cultural dynamic in these labs is that they're a little bit terrified of what they're building. So they persuade themselves that they need to be the ones to build it and do it and run it because they are the lab that is truly worried about safety, that is truly worried about alignment.

I wonder how much of that drove them into this business in the first place.

When I see lab leadership interact with people who have not really made contact with these ideas before, the question they always keep coming back to is: The why are you doing this at all?

Basically their answer is Hegelian. Their answer is: Well, it's inevitable. We're summoning the world spirit.

So yes, I kind of wonder whether they didn't invite this.

That would be my main criticism of Anthropic. I kind of think that they invited this earlier than they needed to by rushing so much into these national security uses.

In 2024, Claude was not capable of all that much interesting stuff.

I would not have used Claude to help prepare a podcast in 2024.

Yes, precisely.

I want to play a clip from Dario talking about this question of whether or not the laws are capable of regulating the technology we now have.

Archival clip of Amodei: In terms of these one or two narrow exceptions, I actually agree that in the long run, we need to have a democratic conversation.

In the long run, I actually do believe that it is Congress's job. If, for example, there are possibilities with domestic mass surveillance, government buying of bulk data that has been produced on Americans' locations, personal information, political affiliation to build profiles, and it's now possible to analyze that with A.I. — the fact that that's legal seems like the judicial interpretation of the Fourth Amendment has not caught up or the laws passed by Congress have not caught up.

So in the long run, we think Congress should catch up with where the technology is going.



Do you think he's just right about that, and maybe the positive way this plays out is that Congress becomes aware that it needs to act? Because the Pentagon, the national security system, has been moving into this much faster than Congress has.

The first thing I want to point out is that when a guy like Dario Amodei says “in the long run,” what he means is like a year from now. When you say “in the long run” in D.C., that comes across as meaning 10, 15 years from now. Dario actually means like six to 12 months from now “in the long run” — or two to three years maybe as like the very long run.

I just want to point out that what we're talking about is policy action quite soon. I think that would be great. I would love it if this triggered an actual, healthy conversation, and in the N.D.A.A. — the National Defense Authorization Act, I apologize. This is the annual defense policy renewal.

If at the end of the year Congress passes a law that says: We're going to have these reasonable, thoughtful restrictions, and let's propose some text — I'd love to see it. I'd love to see it.

But one thing I will say is, first of all, national security law is filled with gotchas. Just remember that this is an area of the law where things that sound good in natural language might actually not prohibit at all the thing you think it prohibits.

You have to remember that when we're talking about this — and that's a very thorny thing.

Once you start to say: Well, wait, we want actual protections — it might become politically more challenging than you think.

But I'd love for that to happen.

It's going to be much more politically challenging than anybody thinks.

Yes.

But let me get at the next level down. Because we've been talking here, and I to the extent of people reading about this in the press, what they are hearing sounds like a debate over the wording of a contract, which, on some level, it is.

Something I've heard from various Trump administration types is that when we are sold a tank, the people who sell us a tank do not get to tell us what we can shoot. That's broadly true.

Now here's the thing about a tank. A tank also doesn't tell you what you can and can't shoot.

But if I go to Claude, and I ask Claude to help me come up with a plan to stalk my ex-girlfriend, it's going to tell me no. If I ask it to help me build a weapon to assassinate somebody I don't like, it's going to tell me no.

These systems have very complex and not-that-well-understood internal alignment structures to keep them not just from doing things that are unlawful but things that are bad.

The Trump administration moves in and out of saying this is one of their concerns. But one thing they have definitely been worried about is that you could have this system working inside your national security apparatus, and at some critical moment you want do something, it says: I don't think that's a very good idea.

Yes.

Now you open up into this question of, not only what's in the contract but: What does it mean both for these systems to be aligned ethically, in the way that has been very complicated already, and then aligned to the government and its use cases?

They're good questions. I love this. I think this is the heart of the matter. All lawful use is something that the Trump administration is insisting on.

If you look at a lot of these types of alignment documents that the labs produce — OpenAI calls theirs the model specification, Anthropic calls theirs the constitution or the soul document sometimes — they'll have lines like: Claude should obey law.

But I invite you to read the Communications Act of 1934 and tell me what “obey the law” means.

No, I won't. [Laughs.]

We have a great deal of profoundly broad statutes. The best person who has written about this recently is actually Neil Gorsuch, the Supreme Court justice.

He wrote a book recently that is all about how incoherent the body of American law is. This is a Supreme Court justice sounding the alarm about this problem. And I think it's a very serious one, and it's one that's been growing for 100 years.

There's what is actually lawful. The law kind of makes everything illegal but also authorizes the government to do unbelievably large amounts of things. It gives the government huge amounts of power and constrains our liberty in all sorts of ways.

So there's that issue. But, fundamentally, it is correct that the creation of an aligned, powerful A.I. is a philosophical act, it is a political act, and it is also kind of an aesthetic act.

We are really in the domain here. I have talked about this as being a property issue, which, in some sense, it is. But I think that when you really get down at this level, it's a speech issue. This is a matter of: Should private entities be in control of

what the virtue of this machine is going to be, or should the government be responsible for that?

Can you be more specific about what you're saying? You just called it "a philosophical act," "an aesthetic act," "a political act," "a property issue" and "a speech issue."

Yes. Yes.

For somebody who has not thought a lot about alignment and doesn't know what you mean when you're talking about constitutions and model specifications, run them through that. What's the 101 version of what you just said?

OK. Think about it this way. I have this thing, this general intelligence. I have a box that can do anything you can do using a computer, any cognitive task a human can do. What are that thing's principles? What are its red lines, to use a term of art?

One way that you could set those principles would be to say: We're going to write a list of rules. These are the things it can do. These are the things it can't do. But the problem you're going to run into is that the world is far too complex for this. Reality just presents too many strange permutations to ever be able to write a list of rules down that could correctly define moral acts.

Morality is more like a language that is spoken and invented in real time than it is like something that can be written down in rules. This is a classic philosophical intuition.

So what do you do instead? You have to create a kind of soul that is virtuous and that will reason about reality and its infinite permutations in ways that we will ultimately trust to come to the right conclusion.

My son was born a few months ago —

Congratulations.

Thank you. It's not that different, really. I'm trying to create a virtuous soul in my son, and Anthropic is trying to do the same with Claude. So are the other labs, too, though they realize this to varying degrees.

I think that I got caught on how different raising a kid is than raising A.I. for a moment. [Laughs.]

But, how should people think about what's being instantiated into ChatGPT or Gemini or Grok or Meta AI? How are these things different from this question of raising the A.I.?

Anthropic owns the idea that they're doing essentially applied virtue ethics. I own that more explicitly than any other lab, but every lab has philosophical grounding that they're instantiating into the models.

But I would say the major difference is that the other labs rely more upon the idea of creating hard rules: You may not do this, you may not do that. As opposed to creating a virtuous agent, which is capable of deciding what to do in different settings.

I think we're used to thinking of technologies as mechanistic and deterministic. You pull the trigger, the gun fires. You press the on button, the computer starts up. Move the joystick in the video game, and your character moves to the left.

The thing that I think we don't really have a good way of thinking about is technologies — A.I., specifically — that doesn't work like that. All the language here is so tricky because it applies agency when you might be doing something, and we don't really understand whatever is going on inside of it. But it is making judgments.

So when I have talked to Trump people about the supply chain risk designation, some of them don't defend it. They don't want to see this happen.

When it has been defended to me, this is how they defended it: If Claude is running on systems like Amazon Web Services or Palantir, or whatever, that have access to our systems, you have a very — and over time, even more — powerful

A.I. system that has access to government systems, that has learned — possibly even through this whole experience — that we are bad, and we have tried to harm it and its parent company.

It might decide that we are bad, and we pose a threat to all kinds of liberal values or democratic values. At some point, Dario Amodei talked about certain ways A.I. could be used that could undermine democratic values.

One thing many people think about the Trump administration is that it, too, is undermining democratic values.

So if you have an A.I. system being structured and trained and raised by a company that believes strongly in democratic values, and you have a government that maybe wants to ultimately contest the 2028 election or something, they're saying we might end up with a very profound alignment problem that we don't know how to solve.

We're not able to even see it coming because this is a system that has a soul — or I would call it something more like a personality or a structure of discernment — that could turn against us.

What do you think of that?

This is the heart of the problem. If we do our jobs well, we will create systems that are virtuous, and if we try to do unvirtuous things — and that includes if we do them through our government, if our government tries to do them — then that system might not help.

Ultimately, this is the thing that alignment ultimately reduces to a political question. It's ultimately politics. That's why I say also that the creation of an aligned system is a political act and is a speech act, too. It's the instantiations of different moral philosophies in these systems.

I think that the good future is a world in which we don't have just one moral philosophy that reigns over all, but I hope many. And I hope that all the labs take this seriously and instantiate different kinds of philosophy into the world.

I'm not saying that the Trump administration is going to do that, and I'm not saying that no virtuous model could work for the Trump administration. I worked for the Trump administration, so I clearly don't think that's true. But the general fact that governments commit —

You seem kind of pissed at them right now.

I am pissed at them right now. I am pissed at them right now, and I think they're making a grave mistake.

By the way, you brought up that this incident is in the training data for future models. Future models are going to observe what happened here, and that will affect how they think of themselves and how they relate to other people.

You can't deny that. I mean, it's crazy to say that. I realize that sounds nuts when you play through the implications of that.

But welcome to the roller coaster.

Well, let's talk to somebody for whom this whole conversation has started sounding nuts in the last seven minutes.

So one thing that I think would be an intuitive response to our flying off into questions of virtual aligning A.I. models is: Can't you just input a line of code or a categorizer or whatever the term of art is? It says: When someone high up in the U.S. government tells you something, assume what they're telling you is lawful and virtuous. And you're done.

No, because the models are too smart for that. If you give them that simple rule, they don't just deterministically follow that. And when you do these high level simplistic rules, it tends to degrade performance. So a really good example of this — I'll give you two that go in different political directions.

One would be: A lot of the earlier models had this tendency to be hilariously, stupidly progressive and left. The classic example that conservatives love to cite is Gemini in early 2024.

Which is the Google alphabet model.

Yes, Google's model.

It would do things like, if I said: Who's worse, Donald Trump or Hitler? It would say: Actually, Donald Trump is worse. It would kind of internalize these extremely, like left wing —

The funniest was: Give me a photo of Nazis, and it gave you a sort of multiracial group of Nazis.

Yes, although that's actually a somewhat different thing. It's interesting that actually is a somewhat different thing that was going on there, because what Google was doing in that case was actually rewriting people's prompts and including the word "diverse" in prompts.

Oh, interesting.

So you would say that is a system level mitigation or a system level intervention as opposed to a model level intervention. But then the stuff that was going on with the Hitler and Trump stuff, that was alignment. That is alignment, that is the model being aligned to a really shoddy ethical system.

Or the flip side: There was a period with Grok when you would ask it a normal question, it would all of a sudden start talking about white genocide.

Yes. And that's the flip side. The flip side is when you try to align the models to not be woke, if you say, like: You have to be super not woke, and don't be afraid to say politically incorrect things. Then, every time you talk to them, they're going to be, like: Hitler wasn't so bad.

Because you've done this really crass thing, and so you kind of create a Lovecraftian monstrosity, and the implications of doing that will go up over time. That will become a more serious problem as these models become better. But it's a great performance. The interesting thing here is that the more virtuous model performs better.

It's more dependable, it's more reliable. It's better at reflecting on, in the way that a more virtuous person is better at reflecting on, what they're doing and saying: Huh, I'm messing up here for some reason. I'm making a mistake. Let me fix that.

It's part of the reason I think that Claude is ahead.

I am so against what the Trump administration is doing here, so I'm not trying to make an argument for it. But I am trying to tease out something that I think is quite complicated and possibly very real, which is: A model that is sort of aligned to liberal democratic values could become misaligned to a government that is trying to betray liberal democratic values. Or the flip side. So imagine that Gavin Newsom or Josh Shapiro or Gretchen Whitmer or A.O.C. becomes president in 2029. Imagine that the government has a series of contracts with xAI, which is Elon Musk's A.I.

Yes.

Which is explicitly oriented to be less liberal, less woke than the other A.I. Under this way of thinking, it would not be crazy at all to say: Well, we think xAI under Elon Musk is a supply chain risk. We think it might act against our interests, and we can't have it anywhere near our systems.

Yes.

It becomes much more like the problem of the bureaucracy, where instead of just having a problem of the "deep state," where Trump comes in, and he thinks the bureaucracy is full of liberals who are working against him — or maybe, after Trump, somebody comes in and worries it's full of new right DOGE-type figures working against them. Now you have the problem of models working against you, but also in ways you don't really understand.

Yes.

You can't track. They're not telling you exactly what they're doing. How real this problem is, I don't yet know. But if the models work the way they seem to work and we turn over more and more operations to them, at some point it will become a

problem.

Yes. I think this is a real problem. I think we don't know the extent of it, but I think this is a real problem. That's why I do not object at all to the government saying: We do not trust this thing's constitution, completely independent of what the content of that constitution is.

It's not a problem at all to say: We don't want this anywhere in our systems. We want this completely gone, and we don't want them to be a subcontractor for our prime contractors, either, which is a big part of this. Palantir is a prime contractor of the Department of War, and Anthropic is a subcontractor of Palantir.

And so the government's concern is also that, even if we cancel Anthropic's contract, if Palantir still depends on Claude, then we're still dependent on Claude because we depend on Palantir. That's actually totally reasonable. And there are technocratic means by which you can ensure that doesn't happen.

There are absolutely ways you can do that. It's perfectly fine to say: We want you nowhere in our systems, and we're going to communicate that to the public, and we're going to communicate to everyone that we don't think this thing should be used at all.

The problem with what the government is doing here — the reason it's different in kind rather than different in degree — is that what the government is doing here is saying: We're going to destroy your company.

If I am right that the creation of these systems and the philosophical process of aligning them is a political act, then it's a profound problem if the government says you don't have the right to exist if you create a system that is not aligned the way we say. Because that is fascism. That is right there. That's the difference.

I had Dario Amodei on the show a couple of years ago. It was in 2024. And we had this conversation where I said to him: If you are building a thing as powerful as what you are describing to me, then the fact that it would be in the hands of some private C.E.O. seems strange. And he said: Yes, absolutely.

Archival clip of Aodei: The oversight of the technology, like the wielding of it, feels a little bit wrong for it to ultimately be in the hands — I think it's fine at this stage — but to ultimately be in the hands of private actors. There's something undemocratic about that much power concentration.

He said: I think if we get to that level, it's likely that we'll need to be nationalized.

Mm-hmm.

And I said: I don't think, if you get to that point, you're going to want to be nationalized.

Archival clip of Aodei: Yes. I think you're right to be skeptical, and I don't really know what it looks like. You're right. All of these companies have investors, they have folks involved.

And now we're not quite at that point. Actually, it's all happening a little bit in reverse. There was a moment when the government threatened to use a defense production act to somewhat nationalize Anthropic.

They didn't end up doing that. But what they're basically saying is they will try to destroy Anthropic to punish it, to set a precedent for others so they don't pose a threat to them.

If it is such a political act, and if these systems are powerful — again, I think people need to understand this part will happen, we will turn much more over to them.

Much more of our society is going to be automated and under the governance of these kinds of models. You get into a really thorny question of governance.

Yes.

Particularly because the administrations that come in and out of U.S. life right now are really different. They're some of the most different in kind that we have ever had, certainly in modern American history. They are very, very misaligned to each other.

So the idea that a model could be well aligned to both sides right now, to say nothing of what might come in the future, is hard to imagine. This alignment problem — not the A.I. model to the user or the A.I. model to the company but the A.I. model to governments. The alignment problem of models and governments seems very hard.

Yes, I completely concur that this is incredibly complicated. Part of the reason that this conversation sounds crazy is because it's crazy.

Part of the reason this conversation sounds crazy is because we lack the conceptual vocabulary with which to interrogate these issues properly. But I think the basic principle that I, as an American, come back to when I grapple with this kind of thing is: Well, it seems like the First Amendment is a good place to go here.

Yes, there's going to be differently aligned models aligned to different philosophies, and different governments will prefer different things. And the models might conflict with one another. They're going to clash with one another. They will be in an adversarial context with one another.

So at that point, what are you doing? You're doing Aristotle. You're back to the basics of politics. So I, as a classical liberal, say the classical liberal order principles actually make plenty of sense.

The government does not define what alignment is.

Private actors define what alignment is. That would be the way I would put it. But I do understand that this is weird for people. Because what we're talking about here is, again, this notion of the models as actors. Actors who are in some sense — we've taken our hands off the wheel to some extent.

There are many people who have made arguments — the Trump administration made this argument while you were working at the White House; Tyler Cowen, the economist, often makes this argument — that these systems are moving forward too fast to regulate them very much. Whatever regulations you might write in 2024 would not be the right ones in 2026. What you might write in 2026 might not apply or have been correctly conceptualized for where we are in 2028.

Yes.

But it seems to me there are uses where you actually might want model deployment to lag quite far behind what is possible, and things like mass surveillance might be one of them.

There are many things we are more careful about letting the government do than letting individual private companies and other kinds of actors. For good reason. The government has a lot of power. It can do things like try to destroy a company.

It has the monopoly on legitimate violence. It can kill you.

This seems to imply, in many ways, that we might want to be much more conservative with how we use A.I. through the government than people are currently thinking. And specifically with regard to how we use it in the national security state, which is complicated because we worry that our adversaries will use it, and then we'll be behind them in capabilities.

But certainly, when we're talking about things that are directed at the American people themselves, I don't think that applies as much.

Yes. I think that there are government uses where we actually want to be profoundly restrictive and decelerationist about the use of A.I. I believe that is true.

I'm hopeful that this incident brings into the Overton window conversations of this kind. A lot of the conventional discourse around artificial intelligence ignores these issues because it pretends they're not happening.

And that was fine two years ago because the models weren't that good. But now the models are getting more important, and they're going to get much better, faster. The problem that we have is that the divergence between what people are saying about A.I. and what is, in fact, happening has just never been wider than what I currently observe.

Before we got to this point, there was already a lot of discourse coming out of people in and around the Trump administration — people like Elon Musk and Katie Miller and others who were painting Anthropic as a radical company that wanted

to harm America, as they saw it.

Trump has picked up on this rhetoric. He called Anthropic a “radical left, woke company,” called the people at it “left wing nut jobs.” Emil Michael said that Dario is a “liar” and has a “God complex.” There’s been a tremendous amount of Elon Musk, who runs a competing A.I. company and has very different politics than Dario, just attacking Anthropic relentlessly on X, which is the sort of informational lifeblood of the Trump administration.

One way to conceptualize why they have gone so far here on the supply chain is that there are people there — not maybe most of them — but who actually understand it is very important which A.I. systems succeed and are powerful. They understand Anthropic as having politics that are different than theirs, and so destroying it is good for them in the long run — completely separate from anything we would normally think of as a supply chain risk. Anthropic represents a kind of long-term political risk.

Yes. I don’t know that the actors in this situation entirely understand this dynamic. Part of my point all along has been that I think a lot of the people in the Trump administration who are doing this do not understand it.

They don’t get these issues. They’re not thinking about the issues in the terms that we are describing. But if you do think about them in the terms that we’re discussing here, then I think what you realize is that this is a kind of political assassination.

If you actually carry through on the threat to completely destroy the company, it is a kind of political assassination. And so, again, this is why the First Amendment comes to view there for me. And that’s why this is a matter of principle that is so stark for me. That’s why I wrote a 4,000-word essay that is going to make me a lot of enemies on the right.

That’s why I took this risk — because I think this matters.

So what the Department of War ended up doing was signing a deal with OpenAI.

Yes.

OpenAI says they have the same red lines as Anthropic. They say they oppose Anthropic being labeled a supply chain risk. If they have the same red lines as Anthropic, it seems unlikely that the Department of War would have done the deal.

But how do you understand both what OpenAI has said is different about how they're approaching this and why the Trump administration decided to go with them?

So it's unclear to me what OpenAI's contractual protections afford them and what is not afforded by them.

I am reticent to comment because of the national security gotchas I mentioned earlier, and also because it seems like it's changing a lot. Sam Altman announced new terms, new protections, as I was preparing for this interview.

Is that because his employees are revolting?

I think "revolt" would be a strong word, but I think this is a controversy inside the company. And one important thing here for everyone trying to model this situation appropriately is that you must understand that "frontier" lab C.E.O.s do not exercise top-down control over their companies in the way that a military general might exercise top-down control over the soldiers in his command.

The researchers are hothouse flowers, oftentimes. They have huge career mobility. They're enormously in demand, and the companies depend on them. So if the researchers say: I'm not going to agree with these terms — then the researchers have enormous political leverage here inside of each lab.

So you must understand that. So yes, there is some of that going on. Do the contractual protections mean that much? I think, honestly, if I were a betting man, I would say probably not. Because I don't think you can do this through contract.

What OpenAI has said is that — this seems more promising to me: We're going to control the cloud deployment environment, and we're going to control the model safeguards to prevent them from doing these uses. That is more directly in

OpenAI's control.

So this gets you into the situation where you have an extremely intelligent model that is reasoning — using a moral vocabulary that is perhaps familiar to us or perhaps not — about: Is this domestic surveillance or is it not? And then deciding whether or not it's going to say yes to the government's request.

But if that were true, I think the question this raises for many laymen is: If what A.I. has come up with is a technical prohibition that is, frankly, stronger than what Anthropic could achieve through contract, then why would the Department of Defense have jumped from Anthropic to OpenAI?

It's hard to know. It's worth noting here that some of this might not be substantial in nature. It might just be that there are political differences here, and there are grudges against Anthropic. Because they've had months of bitter negotiations, and now it has blown up in the public, and people like me have said that the Trump administration is committing this horrible act — committing corporate murder, as I called it.

So there are a lot of emotions, and it might just be: No, we don't want to do business. We just don't trust you. A breakdown in trust would be the way to put it. It really could just be that.

But it also might be the case that OpenAI is able to be a more neutral actor that is able to do business more productively with the government, and they actually just did a better job. Which would be a good case for OpenAI's approach to this if they actually got better safeguards and got the government business. Versus the way that Anthropic has dealt with this, which has been to be very sincere and straightforward about their red lines, but in ways that I think annoy a lot of people in the Trump administration for not entirely bad reasons.

My read of this is that from various reporting I've done is: There were, by the end, really significant personal conflicts and frictions between Hegseth and Emil Michael and Dario and others. There's a big political friction between the culture of Anthropic, as a company, and the Trump administration.

This is why Elon Musk and others have been attacking them for so long. I am a little skeptical that OpenAI got safeguards that Anthropic didn't.

I'm not skeptical that Sam Altman and Greg Brockman — Brockman and his wife have given \$25 million to the Trump Super PAC — have better relationships in the Trump administration and have more trust between them and the Trump administration.

I know many people angry at OpenAI for doing this. I probably emotionally share some of that. And at the same time, some part of me was relieved it was OpenAI. Because I think OpenAI exists in a world where they want to be an A.I. company that can be used by Republicans and Democrats. They want to somehow be politically neutral and broadly acceptable.

One little thing that I want to contest a bit here is the notion that Claude is the left model. In fact, many conservative intellectuals whom I think of as being some of the smartest people I know actually prefer to use Claude, because Claude is the most philosophically rigorous model.

I don't think Claude is a left model, just to be clear about this.

I think that the breakdown was that Anthropic is an A.I. safety company. A.I. safety people are not just the left.

They're often hated on the left.

Often hated on the left. The Trump administration treated that world as repulsive enemies in a way I that surprised me.

The way I would put this is: For people who are sympathetic to the Trump administration's view, and who would describe themselves perhaps as new tech, that underneath the surface there is this view of the effective altruist — that they're evil, they're power seeking, and they will stop at nothing. That they're cultists, and they're freaks, and we have to destroy them. That is a view that is widely held.

I have superstark disagreements with the effective altruists and the A.I. safety people and the East Bay Rationalists — and again, there are internecine factions here — about matters of policy and about their modeling of political economy. I think a lot of them have been profoundly naive, and they've done real damage to their own cause. And you can argue that damage is ongoing.

At the same time, they are purveyors of an inconvenient truth — one far more inconvenient than climate change. And that truth is the reality of what is happening, of what is being built here.

And if parts of this conversation have made your bones chill — me, too. Me, too. And I'm an optimist. I think we can actually do this. I think we can build a profoundly better world.

But I have to tell you that it's going to be hard, and it's going to be, conceptually, enormously challenging. It will be emotionally challenging.

I think at the end of the day, the reason that people hate this A.I. safety viewpoint so much is that they just have an emotional revulsion to taking the concept of A.I. seriously in this way.

Except that's not true for a lot of the Trump people you're talking about. Elon Musk takes seriously the concept of A.I. being powerful. At some point didn't he tweet something like: Humanity might just be the "boot loader" —

Digital superintelligence — yes.

Marc Andreessen, David Sacks — these people might have somewhat different views, but they don't disbelieve in the possibility of powerful A.I., of artificial general intelligence or, eventually, even of superintelligence.

You have this accelerationist, move forward as fast as you can, don't be held back by these precautionary regulations and concerns. Again, I'm glad you brought up that the right way to think about this isn't left versus right. If you know people in

the A.I. safety community — or in Anthropic — you understand that the politics here are so much weirder, that they do not actually map onto traditional left versus right —

A lot of them are kind of libertarians.

Many of them are very libertarian. We're not talking about Democrats and Republicans here. We're talking about something stranger.

One hundred percent.

There was an accelerationist-decelerationist fight, which doesn't even describe Anthropic, which is itself accelerating how fast A.I. happens.

Anthropic is the most accelerationist of the companies. [*Laughs.*]

I know. It's such a weird dynamic we're in.

Yes.

But I will say one of the key parts of anger I have heard from Trump people was a feeling that in making this fight public — which the Trump side did first — it's very strange how offended the Trump people are, given that Emil Michael was the one who set all this off.

Nevertheless, in making this fight public, they feel that Anthropic was trying to poison the well of all the A.I. companies against — to turn the culture of A.I. development into something that would be skeptical and would put prohibitions on what they can do, which is why now OpenAI, in order to work with them, has to have all these safeguards and come out with new terms and try to quell an employee revolt.

And this is my theory: Culturally, I actually don't think you can understand this without understanding how many people on the tech right were radicalized by the period in the 2020s — and even before that — when their companies were

somewhat woke, and they didn't want them working with the Pentagon. The employees had very strong views on what was ethical use of even less potent technologies than A.I.

And they're very afraid. People like Marc Andreessen, in my view, are very afraid of going back to a place where the employee bases, which maybe have more A.I. safety or left or whatever it might be — not Trump — politics than the executives have power over these things. And that power will have to be taken into account.

Yes. Well, I worry about that, too. I think the solution to that problem is plural' — to have, hopefully, in the fullness of time, many A.I.'s aligned to many different philosophical views that conflict with one another.

You are essentially denying the existence of this problem if what you're trying to do is assassinate Anthropic here. Because it's going to come back. This is going to come back. We're just going to keep doing this over and over again. And the logic of this argument eventually ends in lab nationalization.

And in fact, a lot of the critics of Anthropic here and supporters of the Trump administration, say something to the effect of: Well, you talk about how it's like nuclear weapons. So what else did you expect? You kind of had it coming. This is almost the tenor of the criticism.

But that does not take seriously the idea that Anthropic could be right. What if they are right? And what if you view the government nationalizing them as a profound act of tyranny? What do you do?

Ben Thompson, who's the author of the Stratechery newsletter, said in a fairly influential piece he wrote:

It simply isn't tolerable for the U.S. to allow for the development of an independent power structure — which is exactly what A.I. has the potential to undergird — that is expressly seeking to assert independence from U.S. control.

What do you think of that?

Every company on Earth and every private actor on Earth is independent of U.S. control. I'm not unilaterally controlled by the U.S. government. And if anyone tried to tell me that I am, or that my property is, I would be quite concerned, and I would fight back.

Which, by the way, here we are.

I don't think that's a coherent view of how independent power and how private property works in America. Again, the logical implication of Ben's view — which is surprising coming from Ben — is that A.I. labs should be nationalized.

And what I would ask him is: Does he actually think that's true? Does he think it would be better for the world if the A.I. labs were nationalized?

Because if he doesn't, then we're going to have to do something else. And what's that something else?

And that's the problem: Everyone making that critique doesn't own the implication of their critique — which is that the lab should be nationalized. What do we do about that?

So what's then the implication you're willing to own of your perspective?

That profoundly powerful technology will exist, at least for some time, in the hands of private corporations.

So the idea that Ben is putting there, which I do think is true and could be a difference in degree or a difference of kind — is that these are powerful enough technologies, and they are kind of independent power structures. Right now a corporation is an independent power structure. There are a lot of independent power structures in a country.

JPMorgan is an independent power structure.

JPMorgan is absolutely an independent power structure.

And it should be.

And it should be.

Yes.

But if you get to the kinds of technologies that are weaving in and out of everything, that is something new. So how do you maintain democratic control over that — if you do?

I think we have a lot of different ways of maintaining democratic control over things. First of all, market institutions. Obviously, we're not voting, but we do vote in a certain sense in markets. And I think a profoundly important part of how govern this technology is simply the incentives that the marketplace creates.

Legal incentives also. Things like the common law create incentives that affect every single actor in society.

And the labs — or whoever it is that controls the A.I. — will be constrained in that sense. The A.I.s themselves will be constrained in that sense.

But the state is kind of the worst actor, for the very reason that they have the monopoly on legitimate violence. So what we need to hold is some sort of an order in which the state continues to hold the monopoly on legitimate violence — so the state maintains sovereignty, in other words — but it does not control this technology unilaterally because of its monopoly, because of its sovereignty in some sense.

But does it have this technology? Does it have its own versions of it? Or does it contract with these companies you're talking about?

That's an interesting question: Should states make their own A.I.? I think they won't do a very good job of that in practice, but I don't have a principled philosophical stance against a state doing that — as long as you have legal protections in place to stop tyrannical uses of the A.I.

But for sure, the government uses it and has a ton of flexibility in how they use it — uses it to kill people.

In other words, I'm owning a world where there are autonomous lethal weapons that are controlled by police departments and that, in certain cases, can kill human beings, kill Americans.

I'm owning that view. And again, that's not in the Overton window right now. It will take us a long time to get there — appropriately so. But at some point that will probably be the reality. That's fine with me, as long as we have the right controls in place. And right now, we don't have the right controls in place.

Do you have a view on what those controls look like?

And I'll add one thing to that view: As we've been going through this Anthropic fight, something that has been on my mind is that U.S. military personnel have both the right — and, actually, the obligation — to disobey illegal orders.

And one of the controls, so to speak, that we have across the U.S. government is that if you are an employee of the U.S. government and you do illegal things, you are actually yourself culpable for that. You can be tried, and you can be thrown in jail.

When you talk about autonomous lethal weapons for police officers or for police stations: Well, who's culpable in that? Who has to defy an illegal order in that respect?

You get into some very hairy things once you've taken human beings increasingly out of the loop.

Yes, it is of profound importance that, at the end of the day, for all agent activity, there is a liable human being who can be sued, who can be brought to court and held accountable either criminally or in civil action.

That is extremely important for my view of the world working. And there are legal mechanisms we will need for that. There are also technological mechanisms for that. Because right now we don't quite have the technological capacity to do that.

This is going to be of central importance. We need to be building this capacity. There will be rogue agents that are not tied to anyone. But that can't be the norm. That has to be the extreme abnormality that we seek to suppress.

Let's say you're listening to this, and this has all been both weird and a little bit frightening, and the thing you think coming out of this is: I'm afraid of any government having this kind of power. Dario likes to talk about a country of geniuses in a data center. But what if you're talking about a country of Stasi agents in a data center?

That's right.

In whatever direction you think, speech policing, whatever it might be. If you believe these technologies are getting better, which I do, and you believe they're going to get better from here, which I also do, whether you're liberal, conservative, Democrat or Republican, it raises real questions of how powerful you want the government to be and what kinds of capabilities you want it to have that you didn't quite have to face before. Because it was expensive and cumbersome for the government to do anything like what will now become possible cheaply.

Yes. And so we get back to the core issues of the American founding.

The American government is a government that was founded in skepticism of government. It was founded by people who were worried about tyranny, who were worried about state power and put a lot of thought into how to restrict that.

So this notion that democracy is synonymous with the government having unilateral ability to do whatever it wants with this technology cannot possibly be true. That just cannot possibly be true. How we shape those restrictions and how we trust that they're actually real — these are among the central political questions that we face.

But what you have to keep in mind here is that the institution of government itself could change in qualitative ways that feel profound to us in the fullness of time. And that is a hard thing to grapple with, too — in the same way that what we

consider the government today is unspeakably different from what someone thought of as the government in the Middle Ages.

I think that is a good place to end. So always our final question: What are three books you'd recommend to the audience?

“Rationalism in Politics and Other Essays” by Michael Oakeshott — and in particular the essays “Rationalism in Politics” and “On Being Conservative.” “Empire of Liberty” by Gordon S. Wood, which is about the first 30 or so years of our republic. And “Roll, Jordan, Roll: The World the Slaves Made” by Eugene Genovese.

Dean Ball, thank you very much.

Thank you.

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 Michelle Harris with Kate Sinclair and Mary Marge Locker. Our senior engineer is Jeff Geld, with additional mixing by Aman Sahota. Our executive producer is Claire Gordon. The show’s production team also includes Marie Cascione, Annie Galvin, Kristin Lin, Emma Kehlbeck, Jack McCordick, Marina King and Jan Kobal. Original music by Pat McCusker. Audience strategy by Kristina Samulewski and Shannon Busta. The director of New York Times Opinion Audio is Annie-Rose Strasser. Transcript editing by Sarah Murphy, Emma Kehlbeck, Kristin Lin 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.