# An Interview with Gregory Allen About Anthropic and the U.S. Government

Thursday, March 5, 2026

Listen to Podcast

**Listen to this post:**

0:00:00 1:16:06

Good morning,

I am happy to welcome **Gregory C. Allen** back for another Stratechery Interview. Allen is a **Senior Advisor** at the Wadhwani AI Center at the Center for Strategic and International Studies, a Washington D.C. non-partisan think tank that was founded in 1962, and the host of the **AI Policy Podcast**. I previously interviewed Allen in **December 2024**, **October 2024**, **October 2023**, **May 2023**, and **December 2022**.

My interviews with Allen have generally

Search...

Stratechery *Plus* UPDATES

An Interview with Gregory Allen About Anthropic and the U.S. Government
Thursday, March 5, 2026

Anthropic's Skyrocketing Revenue, A Contract Compromise?, Nvidia Earnings
Wednesday, March 4, 2026

Technological Scale and Government Control, Paramount Outbids Netflix for Warner Bros.
Tuesday, March 3, 2026

View All

Stratechery *Plus* PODCASTS

How High Can the Hornets Fly?, Book Club on Cavs-Pistons, Mail on Wemby, Desertion Rates, and Coach Flopping
Greatest Of All Talk | Mar 6

centered on the intersection of technology and U.S. defense, and it's hard to imagine a more momentous week to discuss exactly that. We spend nearly the entirety of the interview focused on the dispute between Anthropic and the U.S. government, including a deep dive on the nuclear analogy, a detour into chip controls, an explanation of autonomous weapons, and why the government is dependent on private companies. We then get into questions of power and sovereignty, and why it might be a blessing in disguise that this dispute is happening now. In addition, we discuss Starlink, including Elon Musk's past disputes with the Pentagon, how Starlink is transforming terrestrial weapons, and why private innovation is so important to the U.S.'s national security.

As a reminder, all Stratechery content, including interviews, is available as a podcast; click the link at the top of this email to add Stratechery to your podcast player.

On to the Interview:

## An Interview with Gregory Allen About Anthropic and the U.S. Government

*This interview is lightly edited for clarity.*

**Topics:**

MacBook Neo
Dithering | Mar 6


(Preview) The Anthropic Mess Continues, Frontier AI and the Uncertain Future of Law, Q&A on Netflix, Dating Apps, F1
Sharp Tech | Mar 6


### Stratechery *Plus* INTERVIEWS

An Interview with Gregory Allen About Anthropic and the U.S. Government
Thursday, March 5, 2026

An Interview with Bill Gurley About Runnin' Down a Dream
Thursday, February 26, 2026

An Interview with Matthew Ball About Gaming and the Fight for Attention
Thursday, February 19, 2026

View All

### Sharp Text ARTICLES

The End of the World As We Know It
Friday, March 6, 2026

The NBA's Problems Are Structural, Cultural and Fixable
Friday, February 20, 2026

Takaichi, Tanking and Legalization Lessons
Friday, February 13, 2026

View All

Oppenheimer and Decision-Making | AI vs. Nuclear Weapons | Chip Controls and China | Military AI and Autonomous Weapons | Private Contractors and the Pentagon | Power and Sovereignty | Starlink

## Oppenheimer and Decision-Making

**Gregory Allen, welcome back to Stratechery.**

**Gregory Allen:** Lovely to be back.

**From what I hear you're basically a professional podcaster these days so this should be this should be an easy lift for you.**

**GA:** It's true and it's my shameful secret but I do plan on getting back to writing here in the not too distant future.

**I'm happy to give you a place to drop some takes. This week of all weeks I think is very interesting sort of at a high level. There's topics we're going to talk about that I think are pretty obvious, and I'm interested not just in your takes, but also the D.C. perspective. The San Francisco perspective on a lot of this stuff, I think, is very well represented online. I am neither D.C. nor San Francisco, I'm also now not Taiwan, but I certainly have my own point of view. I think there's something to be said.**


## More by Ben Thompson

### Year in Review

*The most popular and most important posts on Stratechery by year.*

View All

### All Articles

*Explore all free articles on Stratechery.*

View All

### All Content

*Explore all posts on Stratechery.*

View All

**GA:** You're carrying a lot of water for Wisconsin.

**(laughing) Yeah, Wisconsin takes on AI. There's something to be said for just being in a place, drinking the water, feeling what's in the air and so in all these questions, I hope we could get just not your view, but what your view of the way things are going down there would be super valuable.**

**GA:** Sounds great.

**All right, but we're actually going to start with a different city, we're going to start with Los Angeles, California, specifically Hollywood. We're doing something new, this is the first ever movie review interview. I think you tried to assign me homework before this interview, asking me if I'd ever watched Oppenheimer.**

**GA:** That's not true! I assumed you had. I asked you when was the last time you watched it, to see if it was fresh on your mind.

**Oh, that's true. Amazing movie, absolutely loved it. A very interesting movie, I think, given the last week. But when you when you sent that question, where was your mind at?**

**GA:** So I saw Oppenheimer when it came

out, but I watched it, I was at the India AI Summit a couple of weeks ago, and I watched it on the plane ride back, so it was very fresh in my mind.

There's so many clear parallels between the people who are developing ultra powerful technologies of military and national security and global consequence, asking themselves the question, "What is our obligation in this story vis-a-vis the use of that technology? Can we feel like we are participating in an ethical activity if we just delegate decisions about whether or not to use the technology over to the military?". The point basically being that like **this debate that Anthropic's leadership is going through**, that the leadership of many companies are going through, it's been with us for a very, very long time.

What's so interesting is that like the Oppenheimer movie was very favorable to Oppenheimer's guilt, right? That he felt, as he learned more and more about what happened to the people in Hiroshima and Nagasaki, that he kind of wished this had never happened. I kind of have an unpopular opinion on this topic, which is, while it's definitely true that Oppenheimer was horrifically mistreated, I think the U.S. government actually got it right in that story, which is to say they were right to develop atomic weapons, they were right

to use atomic weapons, they were right to develop hydrogen fusion weapons, and they were right to insist that they, the government, be the ones who have control, not the scientists who were making decisions on these things.

**Right, and I appreciated in that movie they included the famous scene of Truman meeting with Oppenheimer and being like, "Don't bring that guy here ever again, this is ridiculous".**

**GA:** Yeah, he's like, "I never want to meet with him again, and shame on you for trying to bring your guilt here, I'm the one who matters here, not you".

**It's interesting. I don't know if you know, I actually have a personal connection to Hiroshima.**

**GA:** Did not know!

**My great-grandma actually died in the bombing, believe it or not.**

**GA:** No kidding.

**But I agree with you. And actually, I think the more you know about the war in the Pacific, which was like my obsession in high school, every single report I could, it was about that. It did save lives. But above and beyond that, I think the point and the other thing that I've gone back**

**to is actually COVID.**

**GA:** Can we say can we say one thing on Hiroshima before we get to the COVID analogy?

**Yeah, for sure.**

**GA:** So I think it is it is widely understood and accepted that the use of nuclear weapons reduced American casualties.

**Oh it reduced far more Japanese casualties.**

**GA:** That is the key. That is the key that I don't think is widely understood. The US Air Force, and even the British Air Force, were in the process of bringing all the bombers that had been bombing Germany to Japan and they were going to be running massive bombing raids effectively 24/7. So even all the casualties that Japan had been experiencing up until that point is actually not a good proxy for just how many people would have died in non-nuclear strategic bombing of Japan under **Curtis LeMay** if the war had continued. So it's very tragic that a lot of Japanese civilians died, but I do think the evidence is actually quite clear that the use of nuclear weapons and the shortening of the war actually saved net Japanese civilian lives.

**It's a great point because those are lives that are saved, that's not talking about**

**the Japanese fighting, it's just talking about lots of civilians die in war, particularly if you're running bombing raids 24/7, it's a great point.**

**But the analogy I wanted to bring up with COVID was one of the things that I found very frustrating about the COVID experience is I think we lost sight of the appropriate place for experts vis-à-vis political decision makers. And of course, experts are important, you want experts on immunology or respiratory diseases and you want economics experts, you want people on supply chains, we certainly could use those before we just shut down shipping for a month. And, you know, to what role did that play in the eventual inflation have used all those sorts of things before we shut down supply chains and potentially contributed to inflation — it's still an open question, not fully resolved — and the point is the experts, the reason they're experts is they are focused on one domain and they are ideally the best at that domain, but political decisions are tradeoffs and tradeoffs imply multiple domains are at play.**

**I think one of the concerns and tragedies of the political class as a whole over the last number of years has been this desire to offload decision-making,**

**and it sounds good to say "We're going to let the experts decide", but when you do that, you're offloading it to folks whose expertise by definition means they're incapable of making tradeoffs because they're specialized, they're not generalized. I just think about this in the context, go back to the Oppenheimer example, to the atomic war example. The reason why it was Truman's decision is because he had the input from the military, he had the input from the various planners, he had the input, the moral input that he had, the input of the scientists and he also was the one that was going to stand for re-election. The democratic point has been brought up in the context of this debate around Anthropic, we can get into that more a little bit, but this is a point that I think is forgotten, is someone has to balance the tradeoffs.**

**GA:** I think you're making two separate points, if I understand you correctly. One is that they are the best positioned person to make the decision. In a sense, they are, "the most likely to get it right", but there's also an additional factor, which is the democratic argument, which is it is just for them to be making the decision.

**That's right. I'm making two arguments, that's exactly right.**

**GA:** Even if they are wrong, even if they are wrong, it should be in their hands.

## AI vs. Nuclear Weapons

**That's exactly right. But I think this analogy, the nuclear weapons and AI analogy, I obviously went straight to it in the Article I wrote this week. What are the similarities? To what extent is that a valid analogy? I'm going to get to the ways that it's not a valid analogy, but let's start out on why it is a valid analogy between nuclear weapons and AI.**

**GA:** I read your piece, of course, and I thought it was quite interesting. And one thing that I do want to tip my hat to **[Anthropic CEO] Dario [Amodei]** and others who are part of that community. I've known about Dario since 2015, and I've known him since I think it was like 2018, I met him early in my national security career. To his credit, he was like one of the genius AI scientists who would show up to meetings when the DoD asked for genius AI scientists to show up to meetings and he did that when he was not rich, when he was not powerful, even when he was at Google Brain before he was at OpenAI. He actually cared, he actually thought about it and him and the community that he was a part of, they took the exponential growth trajectory very seriously in AI performance.

So, you know, the scaling laws, the early

discussions of the scaling laws, they said, "At this point, AI infrastructure spend will be X percentage of GDP" and they also said, "At this stage, AI capabilities will be so significant that conversations about nationalization of the AI labs will have to be taken very seriously". And so they've been anticipating all of this stuff for a very long time, which is not to say that they're always going to make a great decision about it. But I do think they took their own arguments very seriously early to their credit.

There's a few things I want to say in response to your question about, "How should we think about AI?", and I think there's there's two analogies that are reasonable and it kind of parallels the, "Is AI a normal technology or is AI fundamentally different?", that you see in the labor debate, like the economic and labor impacts of AI.

**Absolutely, I was going to make that point later. This specific question actually sort of comes to bear on everything about AI.**

**GA:** So let's first use the normal technology analogy and here I think the best analogy is computers. So if you go back to 1945, there's like one, I think it was still one in 1945, one digital programmable computer on planet Earth. And it has one use, which is code breaking. That was the original killer

app of computers.

**Was it ENIAC?**

**GA:** No, ENIAC came later. I'm talking about **Colossus**, the UK one for breaking the Enigma codes.

**Oh, Enigma! I knew there was an E involved there somewhere.**

**GA:** So the point basically being that computers started as an application-specific technology, but obviously they became a general purpose technology, and so as you fast forward 80-plus years later, computers are not just used for code-breaking, they're used in airplanes, missiles, ships, tanks, and even the stuff that the military does that doesn't have computers inside it, like bullets or clothing was designed on a computer and the logistics that got it to warfighters was managed on computers.

So that's the first analogy where AI could play out is it's like computers, which is to say, when I was in the Department of Defense, we had a few use cases where AI was immediately delivering really impactful value. Then slowly but surely, the use cases of AI keep diversifying such that it eventually, in a decade or so, we expect it to touch everything that the DoD does, but it's everything the DoD does in the way that computers did. So that's the first scenario

for AI.

The second scenario for AI is the nuclear weapons analogy, and you can think about all of human history from a geopolitics, the history of geopolitics standpoint, as like being divisible into two eras, the before nuclear weapons era and the after nuclear weapons era. Like literally everything is different once nuclear weapons exist on the table in international relations.

**Right. The entire Cold War concept is totally bizarre to anyone who does not understand what nuclear weapons are and why they exist.**

**GA:** Yeah, like there have been strategic standoffs that did not go hot before nuclear weapons, but the specific architecture and incentive structure really does depend upon nuclear weapons and the way that it manifested in the Cold War.

So how how might AI be like nuclear weapons? Well think about for example you know the exponential growth of **Moore's Law**, which since the 1960s has seen performance-per-cost of computers go 2-4-8-16-32 such that you know the computers we have right now are literally a hundred billion times better than the stuff that **Gordon Moore** was writing about in the first data point on his chart, so we're used to crazy improvements in

performance.

But AI, you can actually point to like four different inputs that are plausibly undergoing Moore's Law-like phenomenon. Not just that the chips are still getting better, and if you believe Nvidia's marketing materials, the stuff that they have right now is like over a thousand times better than the stuff that they were selling less than 10 years ago. Then there's the number of chips. I think GPT-3 was reported as training on like 10,000 or so V100s and now we have clusters that have millions of Blackwell's in them, so those are two things that are increasing at exponential pace. The next one is the amount of training data that is used, which even though the Internet is no longer this amazing fossil fuel type resource, reinforcement learning is still generating insane increases in data and the hooking up of AI systems to sensors is still generating an incredible amount of new data. And then the fourth trend is the algorithmic efficiency, like how much intelligence juice can we squeeze out of a given input of computational resources and data resources.

All four of those things are undergoing exponential performance increases and what this means is that instead of AI performance going 2-4-8-16-32, it might be going 10-100-1,000-10,000, and so what

that means is that even if you use Claude, even if you use ChatGPT and you think, "This sucks, it can't do this part of my job, it can't do this part of my job, can't do this part of my job", we might be very far away in terms of performance from strategically transformative capabilities, but we might not be that far away in terms of years from strategically transformative capabilities.

Let me give you a national security example. Right now, or actually maybe not right now, but the example I used a year ago was that the NSA, the National Security Agency, which is the part of the U.S. intelligence community and the Department of Defense that has like the most and best hackers, has been labor-constrained for a really long time. Say they have I think their total employee count was something like 15,000, don't quote me on that, but not all of those are awesome hackers. So let's say the awesome hacker workforce is on the optimistic side, 5,000 people. Well, it takes a decade to make an awesome hacker at the NSA, and you can't just magically go from 5,000 of these people to 500,000 of these people unless AI agents are now the hackers, and so you could imagine a scenario where if we do have asymmetric — as in we have it for some period of time and China doesn't — AGI or ASI type capabilities, we could go from 5,000 hackers to 5 million AI hackers,

each of whom is as good as the best folks in the NSA. That means you could literally devote like an NSA quality strike team to every single member of the People's Liberation Army in a simultaneous first strike and that is not some kind of productivity enhancer, that is a strategically geopolitically relevant new capability that would exist.

And again, this is just one example, but the Cold War, the way the **game theory** ended was **mutually assured destruction**. But one of the fathers of game theory, **John von Neumann**, that was not his preferred outcome of the Cold War. Do you know what his preferred outcome was?

**We use nuclear weapons, we bomb them.**

First strike while they don't have nukes. And this guy was definitely advising the most powerful people in the country at the time. He's saying, "What are you all doing?". If you ask me the question, "Why should we not nuke Russia tomorrow?", I would respond, "We shouldn't, we should nuke them today", that was what John von Neumann said about the game theory outcome of nuclear war.

And so, again, when it comes to geopolitics, even if you think everything I've said is crazy, the geopolitical implications do not

require me to be correct, they just require powerful world leaders to think I'm correct, the truth separate from the belief as variables in the equation. And so those are kind of the two scenarios for AI's impact, the computer analogy and the nuclear weapons analogy.

**Well, on that point, and this is something that I've sort of brought up continuously, and I think I've had frustration, frankly, with Dario specifically, where he talks regularly about it's important to not sell ships to China and that's where he often brings up the nuclear weapon analogy. And the reason I'm frustrated about that is in those comments, he never mentions Taiwan. And again, I'm biased on this point because I used to live there, but my point is, if we actually achieve this asymmetric capability, the game theory doesn't stop with AI. There's still guns and missiles, and if we have that capability, it seems at least a reasonable assertion or reasonable worth taking seriously that the game theory optimal play for China is just to blow up TSMC. If you're going to have asymmetric AI capabilities, then no one gets AI, period. And as long as that's the case, that seems to be a factor that should weigh on these decisions.**

**GA:** So I certainly agree it's a factor. I don't necessarily agree with you that the implications of that thought experiment are as obvious as you suggest. To me, there's still additional variables in the equation. Number one is if we do have something approaching AGI or ASI, how long do we have a monopoly? How long can we keep our monopoly a secret? And then how potent are first strike capabilities and how rapid do those first strike capabilities manifest themselves? If we could say, "We have enough hackers that we could plausibly disable their missile force". And just so everybody knows, this has been in the New York Times, people try to hack each other's nuclear missiles, that's definitely a thing that militaries do.

**So do you think that is their angle? That actually I'm underestimating the extent to which ASI can make this a moot point.**

**GA:** I think if you're trying to model Dario's brain, asking yourself, "Am I underestimating how AGI-pilled he is?", is always a good heuristic.

**(laughing) Good point. I mean, it's funny. This came up in some of the debates I saw about my piece online. I would think, actually, relatively speaking, I'm fairly AGI-pilled.**

**GA:** There's like a whole other level, man.

**Right. No, I know. But it's like the – because a lot of my takes are more about, speaking of game theory, just the competitive dynamics, that we can't actually stop this. So it's less about I think any devotion to that point is missing the point, which I think Anthropic is actually the best example. This is the company that is led by the safety folks who are pushing harder and more aggressively than basically anyone, which I think speaks to the power of capitalism and market forces and whatever it might be. And also, yes, we want to be first, whatever motivations are there, you can't underestimate those to be sure, and given that, we have to assume it's going to happen and then figure out what to do as opposed to try to stop it completely.**

## Chip Controls and China

**GA:** I could launch into my you know my sense of Taiwan and the chip export controls — do you want me to do that now or did you want to go ahead?

**We were going to save for the end here we are.**

**GA:** I think there's a few things here. One thing that's worth noting is that I have — in D.C. you will often see people who are very cruel to the Bill Clinton-era foreign policy to the W. Bush-era foreign policy, talking

about how naive they were vis-à-vis China and basically the theory that engagement would bring peace. And unusually, I'm going to defend that crowd for a moment because there was definitely a time when you looked at the composition of the Chinese Politburo when the, "Yeah, we want to find our way to becoming the next Singapore" of not a full democracy, but like certainly more liberal than our current, semi-totalitarian system and engagement as like a ladder that they would climb towards increasing liberalization. That was not just something that people in Washington, D.C. believed, that was something that powerful people in Beijing believed.

The key here is that Xi Jinping has purged all of those people. There is nobody within like 100 yards of power in China who thinks that anymore, and so when you think about Taiwan, where people often ask, "Why does Taiwan want to buy F-16s? Why does Taiwan want to buy these things?", these are not actually the capabilities that would matter the most to the defense of Taiwan. Well, it's because their military is not actually their national security strategy, their national security strategy is TSMC and remaining essential.

I think, unfortunately, that was a good national security strategy in 2011 and I don't

think it's a good national security strategy in the Xi Jinping era, because I do believe that he is, come what may, willing to do whatever it takes to eliminate Chinese dependence upon TSMC. I actually don't think it's in the option set to preserve long-term Chinese dependence upon TSMC, because I think even if we were willing to eliminate all restrictions, I still think they would be willing to say, "Okay, but every company has to have a plan to transition off of TSMC capabilities, we're going to provide minimum, quotas of what share of your data centers has to be domestic, and that's going to increase to 90 to 100% over time". All the sort of, "TSMC has a natural monopoly", "ASML has a natural monopoly" arguments—

**Those are economics-based arguments and you can be irrational.**

**GA:** And I don't think they wrestle with the reality of a country that has the industrial scale that China does, and the political will that China does, and the intelligence and espionage capabilities that China has. It's not an apples-to-apples comparison, to ask, "Could Japan create an alternative to TSMC and ASML", versus, "Could China create an alternative to TSMC and ASML?".

**I have two responses to that. Number one, do you think that China would have that view absent what ZTE and then to**

**Huawei?**

**GA:** Oh, my gosh, yes. So you and I have had this conversation multiple times over the course of the years and let the record show I was against ZTE when it happened.

**Me too.**

**GA:** This is a one-way door. And so my argument ever since then has been in for a penny, in for a pound. And I think your argument has been maybe it wasn't a one-way door and I'm like, "No, it's still a one-way door".

**I very well might be wrong.**

**Number two is I think I can actually mount a much stronger defense of the Bill Clinton, W. Bush era policy with China. I'm not saying I believe this, but there actually an even stronger defense which is, it worked. The number one reason to believe that the U.S. and China would not go to war and will not go to war and we will somehow muddle through with regards to Taiwan is because the economies are so deeply enmeshed that it actually is impossible, and again I'm not sure I believe that, but that if you really want to make the case and go full-throated in defense of that era, it actually has nothing to do with China liberalizing or not, it's that countries that are this enmeshed don't**

**go to war.**

**Now, you should go back and read pre-World War I history if you want to stand by that but that is sort of the point that you could make.**

**GA:** Yeah, I think what Bill Clinton and George W. Bush have, if you want to give them credit and if you want to make the argument you just made, which is that the policy worked.

**And by the way, I was reminded this week I need to make caveats about what I'm actually saying is good or bad versus just making analysis and putting it out there!**

**GA:** People taking your argument further than you intended on the Internet? Who could have known!

I think the maximum I would be willing to concede that W. Bush and Clinton achieved is that war is a bad idea. But war being an obviously bad idea is no guarantee that war will not happen.

**Right. You can really get yourself in trouble, if you were say you were attacked, would you sort of respond by attacking all your neighbors who are theoretically going to be neutral? Probably not a good idea.**

**GA: Some countries have made that choice**.

**Some companies can be pretty irrational, to say the least.**

## Military AI and Autonomous Weapons

**All right, well, let's circle back to this. I think you just touched on this actually, I thought in an illuminative way with the NSA. But what does the Pentagon use AI for today? And what are the hopes to use it for in the future? Because you think about war, you're thinking about missiles and artillery and things like that. But you just gave a great example of something that is completely digital, which that's where AI is going to have the biggest impact first anyway.**

**GA:** So let me first just remind folks that, you know, when I was in the DoD, AI was not a synonym for LLMs, it was a synonym for machine learning. The use cases of AI for machine learning were still very active then. And to give you an example at the unclassified level because I have had a security clearance and thus have to be careful about not ever revealing classified information, but Maxar, **which is a commercial satellite company**, they take pictures from space. They sell those pictures to hedge funds, but they really got their start selling them to the United States

intelligence community, the origins of the many acquisitions that became what is now Maxar.

That company has a fleet of low Earth orbit satellites and when I was talking to their chief technology officer back in 2022, he pointed out that they take so many pictures that if you were to take like a satellite imagery analyst, which is a reasonably highly trained individual, and ask them to look at all of the pictures and put together analysis, that they collect 85 years worth of full-time work per day. So one way to solve that problem is to hire 85 times 365 times three shifts a day, that's how many analysts you need. But that story doesn't end in being a profitable company, which is what Maxar would like to be.

So they used AI, especially computer vision, as the sort of first pass analyst and basically say, "Hey, this picture is of empty water", "This picture is of empty water", "This picture has an oil tanker", "This picture has a warship", "Human analyst, here are the pictures I recommend you look at first" — that's like where we were in the 2017 to 2022 era it was like those kinds of applications today that stuff has now actually been unified with LLMs. When we talk about multimodal models it's not just that the model might say, "Hey human analyst I think you ought to look at this

first", it might actually write the first draft of the report, which is a pretty cool capability.

And then machine learning AI also opens up the door to new sensor modalities and ways to gather intelligence. So to think about like acoustic intelligence, for example, drones are really small, which means they're not always a great fit for radar-based detection. And by contrast, they often make a lot of noise.

**Yep. Everyone at the Olympics knows that.**

**GA:** Yes. And so Ukraine actually has this network of acoustic sensors that are just listening all the time in the same way that your Alexa in your house is constantly listening all the time.

**It's fascinating, sonar for land. Who knew?**

**GA:** Yes! Sonar for land and air, exactly. And they're basically able to give early warning of drone-based attacks through that kind of thing, so those are sensor use cases of AI.

We already talked a little bit about the cybersecurity, but just to oversimplify the story, if Claude Code is good at code, you know what's made out of code? Cyber weapons and also cyber defenses. You can go through a code base and basically say,

"Hey, what are all the cyber vulnerabilities that are not patched in this code base?", and Anthropic has, in fact, done this in open source repositories. They just go through open source repositories and identify known cybersecurity vulnerabilities and patch them all and just throw them back out into the ether. So AI is useful for upgrading both your cyber offense and your cyber defense and what the net outcome of that is still very much uncertain, like whether it will have a net beneficial impact to defense or offense is still hotly debated in the field.

Okay, then you go to weapons and weapon autonomy. So first is like a sensor part of it and it's providing capabilities that we in many cases already have.

**But we don't have the capability to process the vast amounts of information.**

**GA:** Well no, where it's just expensive to do it. So to give you an example, like a Tomahawk cruise missile is an awesome missile, it can like hit a target to within an accuracy like one to 3 meters from 300 kilometers away, that's awesome. It can even do that when in a GPS-denied environment in some cases. The problem is it costs like $4 million a shot, which stinks. And how does it work when there's no GPS? It actually has a ground-facing radar

and it looks at how high this hill is, how low that valley is, and it has a topographical map of the entire earth and it basically says, "Okay, if there's this elevation change profile, I must be over this part of the Earth", which is a super cool capability that we like literally figured out in the 1980s with handcrafted algorithms, well before machine learning was a part of the story.

So we would love to have that kind of capability, but not have it cost $4 million a shot, and it turns out that now just using like a webcam and computer vision and off-the-shelf AI algorithms, you can do a cheap and crummy replica of that capability where the drone looks at the ground and says, "If the ground looks like this, then I must be here". So that can get you a long-range precision strike capability that can work in GPS-denied environments that doesn't cost $4 million a shot. Maybe it costs $1,000 a shot or $20,000 a shot or $100,000 a shot and so the the value of AI is not always providing some awesome, unprecedented skill. Sometimes it is providing this hugely precedented skill, but making it much cheaper and more available in classic **Clayton Christensen** disruptive innovation ways.

**And then to your point, that making it cheaper is sort of analogous to the first pass of, "We can get rid of all the**

**satellite images of water", but then the next one is, well, the second one is, we can actually not just say, "Oh, this image has a ship in it", it can identify the ship, it can write a report about the ship and actually do the job.**

**GA:** And connect the dots. Connect the dots to other sources of intelligence.

**The analogy here with weapons is, "I'm flying over a tank, I recognize that's a tank, I'm going to now attack the tank".**

**GA:** So what you've gotten to is the autonomous weapons part of the story, which has like dominated the debate forever. I wrote in 2017, **one of the first major reports** that was like an analysis of AI and national security, but that was just because I was like the first broad survey. There was still a deep literature about autonomous weapons that had long preceded that so it's probably worth just like recounting a few definitions here.

There is no definition of an autonomous weapon in international law. The United Nations had a group whose job was to decide the definition, they spent eight years, they did not agree upon a definition at the end of that group.

**Which sort of makes the point.**

**GA:** Yes. But there is a definition in DoD

policy, and it's in **Department of Defense Directive 3000.09**, and this is what's surprising to a lot of folks. But the definition of an autonomous weapon is one that can, once activated, without further human intervention, select and engage targets under its own autonomous activities.

So the key here is that AI used to identify a target is not autonomous, it has to select and engage the target. So to give you like a stupid example, let's say you had a drone and it had a facial recognition capability and you said, "Drone, your job is to kill Vladimir Putin, here is what he looks like", and then the drone flies around Moscow looking for faces that match Vladimir Putin, and then it blows up the one that it says that has Vladimir Putin's face. Under DoD policy, that is not an autonomous weapon. An autonomous weapon would be one that you say, "Drone, go fly over Moscow and find good targets and when you see a good target go kill it".

What's interesting is that autonomous weapons exist in the arsenals of dozens of militaries around the world and have existed in the U.S. military for decades. The Patriot missile defense system for example has a fully autonomous mode, the Aegis weapon system for example has a fully autonomous mode, so autonomous

weapons are not some new ground in military technology. What is kind of new is the use of machine learning to enable those autonomous functionalities and using those functionalities for offensive use cases.

And just in the past 12 months, we have now crossed the threshold where, according to Ukrainian intelligence that I do find credible, Russia is now deploying AI-enabled, offensive, lethal autonomous weapons that would meet that definition under U.S. policy, which is to say you let the weapon loose and you don't know what it's going to attack. And if it sees something that weapon, even if it loses backhaul communications, it's not being remotely piloted it makes the decision, "That's a worthwhile target, I'm going to go kill it", and it never asks a human for permission.

So autonomous weapons exist on the battlefield today using AI and that's why this Anthropic spat with the military, you can imagine why the military is sort of pulling their hair out, they're like, "Why would we give a company the right to tell us whether or not we can match a capability that Russia has right now?".

**This gets to, I think this is a perfect example of the Truman politician point that we were talking about before. We just had three U.S. F-15s shot down by**

**friendly fire, and the point is mistakes happen in a theater of war. And I do wonder to what extent there is expectations for perfection or the expectations for what is reliable and what isn't not only are going to vary between the military and a private company, they're also going to vary based on circumstance. Like if you — I think this is something that you saw some anecdotes come up — but if it is a crisis situation, your relative balance of what is a risk worth taking is going to change in real time.**

**GA:** In the Department of Defense, there's this system safety review policy called **Military Standard 882 Echo**. And it basically says, if you're going to deploy it in casual circumstances, that's fine, but if it's like potential friendly fire, potential loss of life, potential use of force, you must achieve reliabilities of blah, blah, blah, blah, blah, and that's the standard. But at the end of that, it just says, "But the commander can accept the risk".

That is the human judgment at the end of the story and to give you an example, when the American military rushed in technology that was "not ready" and it really mattered would be the Korean War. So there was a period of time in the Korean War where the Soviets, who even though they were "not in

the war", actually had Soviet pilots flying Soviet aircraft killing Americans and everybody just agreed to keep it a secret.

**We were flying P-51s and they were flying jets.**

**GA:** Exactly. So we had propeller-driven piston aircraft and the Soviets had jet aircraft and we actually rushed into the war these jet fighters that were "not finished" because we said, "Look, we haven't finished all the tests that we would love to run, and so there's a possibility that mechanical failures or other defects are going to lead to a loss of American pilots", but you know what's leading to the loss of a lot of American pilots? Soviet bullets, "So get these things in the field, we accept the risk".

So you can understand how the Department of Defense basically says, "It has to be up to us, we have to be able to accept the risk and decide how stuff is going to be used in operational conditions". I actually have criticism for both Anthropic and for the Department of War in this story. But fundamentally, I'm sympathetic to the Department of War argument that, "We have to be in control".

## Private Contractors and the Pentagon

**Let's go back to go very back to the**

**beginning. We talked about the similarities between nuclear weapons and I think the differences are also very interesting and I wrote about this in a follow up this week. The Manhattan Project cost. How much did it cost? Do you know inflation-adjusted?**

**GA:** I think you said $36 billion recently. I did my own analysis not that long ago, but I don't remember what I found.

**All the news reports say $35 billion, but that's because they need to update for inflation. It is now up to – I believe it's $35.8 now, so it's up to $36 billion. This reminds me actually – an anecdote that I bring up again and again is one of our prior interviews where you were like, "Oh, they're committing real money to this", and you're like, "$6 billion", and I'm like, "I'm sorry, was that a B or a T?" and just like you look at this with tech, I mean, just the $700 billion being committed to CapEx this year for AI. But this is the culmination of a story that goes back to the 60s and I think really Bob Noyce, who I credit as being at the forefront of this, really understanding that if we're going to innovate and push forward on chips, we need a much larger market than the U.S. government. And so in the long run, we will sell lots of chips to the U.S. government, not**

**because we took their R&D money, but because we are going to have the best chips and they will just buy them from us off-the-shelf. And famously, the whole 1980s restoration of the U.S. military was basically based on accepting that reality, "We're going to do off-the-shelf".**

**What's interesting about this, I think, and this really highlights the difference, is at the end of the day, General Groves and New Mexico, that was a U.S. project. There wasn't any question as to whether or not Oppenheimer's angst was ultimately totally pointless, he never had control, was never going to have control. But the challenge is as computers and now AI again puts this on steroids, if the necessity is that these things are developed privately, not because we believe in free enterprise, although of course you and I do, but because it's the only economically viable way to do so, that leads directly to these questions. Suddenly you actually do need contracts. And to me, that's like where the analogy falls apart. Beyond all the other things where uranium is trackable and chips versus code and all those sorts of things, it's just the economic circumstances and what that means for where control starts, the path dependency of the control question is**

**just totally different.**

**GA:** The degree of control that Anthropic wanted, I think it's worth pointing out, was comparatively modest and actually less than the DoD agreed to only a handful of months ago. So the Anthropic contract **is from July 2025**, the terms of use distinction that were at dispute in this most recent spat, which was domestic mass surveillance and the operational use of lethal autonomous weapons without human oversight, not develop — Anthropic bid on the contract to develop autonomous weapons, they're totally down with autonomous weapons development, it was simply the operational use of it in the absence of human control. That is actually a subset of the much longer list of stuff that Anthropic said they would refuse to do that the DoD signed in July 2025.

**Right. July 2025, that's the Trump Administration.**

**GA:** That's the Trump Administration, and that's **Undersecretary Michael**, who's been there since I think it was May 2025. And here's the thing, like the DoD did encounter a use case where they're like, "Hey, your Terms of Service say Claude can't be used for this, but we want to do it", and it was offensive cyber use. And you know what happened? Anthropic's like, "Great point, we're going to eliminate that",

so I think the idea that like Anthropic is these super intransigent, crazy people is just not borne out by the evidence.

**Well, there's also the fact that they are the only classified LLM that's available to be used.**

**GA:** They ran to it initially, right? They sprinted to be supporting U.S. national security in the early days and I think I was in the DoD in the aftermath of Project Maven, which in the summer of 2018, when was Google pulled out of **Project Maven**, the program that Anthropic is now the key AI provider to.

**Is that how Anthropic got in? Was it was that downstream from being partners with AWS?**

**GA:** No, the way that Anthropic got in is actually kind of hilarious, totally like a bank shot, **Jack Clark** has disclosed this recently.

It started out with his safety concerns they were basically saying, "Hey we're worried that our capabilities might lower the barriers to entry to developing nuclear weapons and biological weapons, can we work with you to test whether or not that's true?", and the Biden Administration was like, "We are also concerned about that, let's work with you to test your model, so let's get your stuff running on Department of Energy, let's get your stuff running on

DoD networks so we can run some classified tests". But then Anthropic works on classified networks and it takes so long to get anything to work on classified networks that this was an incredible advantage over all of their competitors so this is why that makes sense, it was just kind of hilarious they from that point on this like totally random not trying to make money use case they had the inside track.

**So who's right and who's wrong here? Let's start with that.**

**GA:** My sense here—

**Number one, I want your sense and this is where I really want what you perceive from D.C.**

**GA:** I think one thing that I just have to say as a caveat to everything I'm about to say is like, you remember when Donald Trump tweeted that the CEO of Intel was like a Communist Party Chinese agent, and then they bought 10% of Intel. So my point is that like as crazy and as heated as the rhetoric has been, I actually think we cannot rule out that a deal might still happen between Anthropic and the Department of War.

**The official paperwork has not come out, it was on its face legally incorrect. Like, that was super clear. And, yes, it is, I think, in everyone's interest to solve**

**this.**

**GA:** That's what I want! What I want is the American military to have access to awesome AI capabilities that provide a meaningful national security edge.

**And I think it's fair to say, to your point, that Anthropic has, by-and-large, demonstrated that they want that to be the case as well.**

**GA:** And if you look at the reporting of what's coming out of the Wall Street Journal or the Washington Post, nobody's complaining about how Anthropic's performing in Iran. People like this this capability.

OK, so who's right and who's wrong? I think the Department of War is right to say that they must ultimately have control over the technology and its use in national security contexts. However, you've got to pay for that, right? That has to be in the terms of the contract. What I mean by that is there's this entire spectrum of how the government can work with private industry.

There's something called the arsenal model, which is like the government actually owns the factories and they build the stuff and then they own it.

**This kind of goes back to my point about atomic weapons.**

**GA:** Yeah, like this still happens in ship-building, for example, where the government owns a shipyard and still builds boats, mostly separate from private industry. Then there is the, "Commercial industry builds it, but the government owns it at the end of the day", so to put it like in satellite terms, space and satellite terms, if you look at like the wideband global Satcom constellation, Boeing built those satellites, but then they sold them to the Department of Defense, and now the Department of Defense owns those satellites. That's in contrast to something like SpaceX Starlink, where effectively the US government is buying it on an as a service terms, they're buying the data off of those satellites, they're buying the comms capacity off those satellites, but they don't own it, SpaceX owns it.

And so my point basically being like, if the government has identified this as an area where they need absolute control, the historical precedent is you pay for that when you need absolute control and, by the way, like the idea that that Anthropic's contractual terms are like the worst thing that the government has currently signed up to — not by a wide margin! Traditional DoD contractors are raking the government over the coals over IP terms such as, "Yes we know you paid for all the research and development of that airplane,

but we the company own all the IP and if you want to repair it...".

**How did that happen? That was one of Noyce's objections to government R&D back in the day is that they took all the IP, but that doesn't seem to happen anymore.**

**GA:** Basically you had some sometimes well-intentioned, sometimes not well-intentioned cloaks of, "You know the government should really take more advantage of the free market system, etc.", and really what that is is corporate lobbyists being like, "How can we trick the government into signing horrific contract terms?". This is the problem when you have like national interest going up against fiduciary duty to investors and they're like, "Well, I love my country", but I also love annual returns and hitting my quota.

So yeah, the DoD signs terrible contractual terms that are much more damaging than the limitations that Anthropic is talking about a lot and I don't think they should, I think they should stop doing that. But my basic point is, I do not see a justification for singling out Anthropic in this case.

**Unless you make the case that AI is different.**

**GA:** Well, here's the thing. Even if AI is different, I still think that the Department

of Defense has made a mistake here, which is the gap between the Anthropic position and the DoD position. Number one, the DoD was able to tolerate this gap and more in July 2025, so saying now we need to go supernova and threaten to kill the company strikes me as going pretty fast down that road in a way that was unnecessary.

But the other thing is I still think we might get a deal because I actually think that the gap between the negotiating positions is not that big. The DoD's policy on autonomous weapons, as I said, is that they are not banned, they are subject to additional technical scrutiny, and they are subject to additional procedural scrutiny. And that kicks in in a senior review process, both when you initiate development and also when you initiate fielding. Anthropic has basically said, "We view the technology maturity as enough to justify beginning development, we do not view the technological maturity as enough to justify fielding". And honestly, that might be where the DoD currently agrees is the story! They might just say, "When we ultimately cross that bridge, we're going to have a vote and you're not, but we agree with you that it's not technologically mature and we value your opinion on the maturity of the technology".

## Power and Sovereignty

**Who gives in on this? Is this Anthropic accepting some sort of face saving slight adjustment or or is the DoD going to give in here?**

**GA:** Well, I think that the first rule of politics in the Trump administration is they declare victory regardless. It is in Anthropic's interest for them to be able to declare victory at the end of this story. So even if, hypothetically, the Department of Defense and Anthropic reach a deal such that Anthropic is not designated a supply chain risk and Anthropic does not have the Defense Production Act invoked against it, I think they'll want to sort of say, "Great job, Department of Defense, you really nailed it here".

**So is Anthropoc in some respects actually getting themselves in trouble with this wave of positive PR that is going along with this dispute?**

**GA:** Oh, yeah. Maybe it's good for their Silicon Valley scene, but I don't think it's good for their Department of War relationship. Because, yeah, you don't want to be viewed as scoring points against this Administration, that never turns out well for you. I mean, until and unless there's a new Administration and that's multiple years away. So that's how the Department of War is wrong.

The Department of War, I think, is also wrong in that the supply chain risk designation is just an egregious escalation here that is also not borne out by what that policy is meant to be used when it's when it's legally invoked and I think that Anthropic can sue and would very likely win in court. The issue is that the Trump Administration has pointed out that judicial review takes a long time and you can do a lot of damage before judicial review takes effect and so the fact that Anthropic is right—

**I think this goes back, though, to your initial point, which is, is this a regular technology or is it something different? I think this is the point that I was trying to, again, I wasn't endorsing this, but the point I was trying to make that I felt no one else was, is that if you have this view that at the end of the day, just to take a very crass, simplistic view that ultimately law comes down to who has the guns, that if you think you're building a gun, that is going to be threatening to the people with the guns. And so if AI — yes, of course, it's not today — but if it is a potential source of power five years down the road, 10 years down the road, two years down the road, let's see how these exponential curves sort of go, it's not just that the government will say, "Fine, we're not**

**going to use you with someone else", they will see you as a rival and want to kill you. And again, I think this is all overwrought as so many of the Trump things are, as you sort of put forward, but I think there's a grain of actual, "It's not insane". There's a point here, again, that I think is lost in all these details, but that was the point I was trying to bring up anyway.**

**GA:** I didn't write down the quote, so I'm not going to quote this verbatim. Jack Clark, who's a co-founder of Anthropic and the Head of Policy over there. I remember he said at a speech that I saw in January 2025, he's like, "Companies like mine in a handful of years might have greater power than most nation-states, and we as a democracy should wrestle with what that really means".

Now I want to point out, we started the conversation by pointing out that it could be the case and probably is the case that the politicians are, for better and worse, the best positioned to make the choice, but also it is just that they be the ones that make the choice. And I just want to add here that when we say the politicians have just authority, part of that is because it's constrained authority. They have the power subject to the limits in the U.S. Constitution, they have the power subject

to the balance, the checks and balances of laws passed by Congress, etc. The law for all of its failings, and there are many, is the expression of democratic will through which legitimate authority is derived, and so when we say that the Department of War must have this powers that still is going hand-in-hand with the fact that they have to follow the law and the law in this case does not allow them to invoke that authority cavalierly.

**That is why, actually, I think the best critique of what I was writing about was embedded in my point, which was if this is so powerful, it is going to in a realpolitik way inspire or push the government to extra-legally seize this such that they disregard the law because at the end of the day, they are the law. That, though, undoes the democratic justification, that's why I personally didn't make that case. I linked to Palmer Luckey's argument about that, but I'm not invoking it because the legitimacy of that case falls apart under my framework because we have a government acting above and beyond the law.**

**GA:** You also talked about the difference between international law and domestic law.

**Well, at the end of the day, there is no**

**difference.**

**GA:** There's a huge difference! There's a huge difference because in international law, there is no sovereign, there is no one with a monopoly on the legitimate use of force.

**Right. So the issue of domestic law is if the sovereign is the bad actor, that's the issue, that's what we're barreling towards. And again, this is why I wanted to – I felt no one was talking about this point specifically, I think I at least raised discussion about it, although a lot of the pushback, maybe I did a bad job, missed this was the point, which was if AI is what it is potentially could be, if it is on the level or even beyond that of nuclear weapons, the sovereign is going to act above and beyond the law. And again, that's why I think the democratic argument does fall apart, because you're talking about an actor that has the capability, because they have guns, to go above and beyond the law and I don't want to get there.**

**I think that's why I wanted to bring this up, because I'm actually quite sympathetic. If Anthropic was going to pick two points, I think they picked two good ones. I think the autonomous weapons was actually a bad choice, I think it's a legitimate concern, but you**

**just, I think, with the Korean War analogy is a great one — a company can't be making these risk decisions, that's the military's job, it's ultimately the President's responsibility. The domestic surveillance one, the loopholes are atrocious, we've gone through this with Snowden, computers make all these laws that we had meaningless because so many of those laws assumed on a certain level of friction and time and energy to do this sort of analysis that we can sort of do at scale. It's a totally legitimate concern. But this gets to the trade-off question is, that's really bad, but I would rather not get to where we got to last week.**

**GA:** I'm so with you. Where we got last week was bad, and that's why I fault the Department of War for having this ultimatum with a time stamp on it.

**There's a bit where if it didn't happen last week, it was going to happen sometime in the next 5 to 10 years.**

**GA:** Yeah, they needed to come to a head. But my point is, I think the parties can make a deal, should make a deal and fundamentally, national security benefits from them.

**I put a little bit of an optimism there at the end, mostly about Congress, but I**

**think there's an optimistic take here in general, which is the industry and government comes to a common understanding. And that this, in retrospect, looks like a tiny little skirmish that actually did set the broader framework for how these interactions are going to happen going forward. And if that happens such that when we get to truly capable AI, far above and beyond what is available now, it's not a situation where then we have the debate and actually the AI company is powerful in a way that they're not today and the government, who is making 1% of their revenue, has no economic power over them, has nothing left but sticks, that's the really bad scenario. So there's a bit, and I mean this genuinely, this might end up being, yes, the Trump administration totally overreacted, but maybe an overreaction in 2026 saves a totally appropriate reaction in 2031.**

**GA:** I actually take that as like a very serious scenario here. Where this is a hugely imperfect analogy, but I like it anyway, I personally think that you could make a strong case that the **Cuban Missile Crisis** saved humanity. Which is to say, if you ran the Cold War a thousand times in a simulator, I think human extinction happens in at least 50% of those scenarios and I

think in all the ones where humanity survives, there has to be a terrifying close call. And the question is, is this our terrifying close call that actually gets people having the right conversations, thinking about these things? Obviously, in the time-deliniated circumstances, I don't think we thought about them in the best light, but at least now we're having this conversation and we can.

## Starlink

**The other thing about this, by the way, is coming to these understandings, of course, a lot of people came in here with a partisan lens and, "Oh, you're you're supporting Trump", or, "You're opposed to XYZ", and the point is we have to come to an understanding that extends across Administrations, it can't be dependent on who's in the White House and who isn't.**

**GA:** Now let's think about some of the precedents like that we're setting here and how they might play out over different Administrations. Elon Musk has been loudly attacking Anthropic every step of the way and basically amplifying every negative message you can imagine. Shockingly, xAI now has a contract to be on classified servers.

Well, I just want to say there's a huge gap between having a contract and being

operational and so there still is no substitute for Anthropic for at least months, at least, there's almost no switching costs to, as an individual consumer, to go from Claude to ChatGPT to Gemini. But once you start building all this stuff into mission-critical national security software, there are switching costs, they're non-trivial.

But now if I was Elon Musk, I'd be like thinking back to September 2022 when I turned off Starlink over Ukraine in the middle of a Ukrainian military operation to retake some territory in a way that really, really, really hampered the Ukrainian military's ability to do that and at least **according to the reporting that's available**, did that without consulting the U.S. government right before.

**What happened there and was there a come-to-Jesus moment for Elon there as well?**

**GA:** There's a couple of things that are data points that are known — one is that Vladimir Putin had made a speech in which he said, "Commercial space assets that are providing military relevant capabilities are fair game", so he's sort of implicitly threatening to kill Starlink. And just so you know, Starlink has a lot of little satellites and so it's hard to have as many missiles as there are satellites to take them out, but

one of the ways you could take them all out is nuke space. As a fun fact, a detonated nuclear weapon in space would generate a massive EMP that would like render — yes, that would destroy like a huge number of satellites.

So when Elon Musk was saying, "I'm afraid that Russia might go nuclear", literally the victim of the nuclear attack might be him, and specifically like his space constellation. So he unilaterally, the best reporting on this topic has been done by Reuters, and according to Reuters reporting, he made a unilateral decision to turn off Starlink activity to a active war zone in Ukraine. That was Ukrainian territory that had been seized by Russia, Ukraine was trying to retake that territory when suddenly they lost access to Starlink and Ukraine's use of Starlink capabilities, it is like super fundamental to everything that they are doing from a comms perspective. Ukraine actually has come to be a really impressive software enabled military force with stuff like algorithms, the line is, "Uber for artillery", where you can summon an artillery strike.

**Turns out, making all those AAA games translated well to doing it in real life.**

**GA:** Yes! Making pretend war, surprising synergies with designing stuff for real war. And so Elon Musk made the unilateral

decision to turn that off and did that without asking the U.S. government.

Now, ultimately, what happened is, as I said at the beginning, if the U.S. government wants control, they need to pay for that and according to Quilty Analytics, they then signed a contract with SpaceX for over $500 million to buy Starlink services for the Ukrainian military. And now Elon, I hope this is within the contract terms, doesn't have the authority to turn it off anymore, he ceded those decision rights.

**What is Starshield and how does that fit into this?**

**GA:** It's just the military branding for Starlink capabilities and it's worth pointing out that we know of Starlink in the constellation as a communications platform, SpaceX also has contracts with actors like the National Reconnaissance Office, which is the spy satellite agency for the United States federal government. There's been reporting about how there are infrared-sensing payloads on some of these SpaceX Starlink satellites.

**Are there dedicated Starshield satellites?**

**GA:** I believe so, yes.

**But is it really just a software switch? They're dedicated ones right now, but if**

**we need to use all of them, we can?**

**GA:** That would be my understanding. And so it gets to the rent-versus-own type capability. How much of the Pentagon's contracting with SpaceX is like, "We pay you such and such millions of dollars and you guarantee us such and such amount of bandwidth over such and such regions per year", versus, "We'll pay for one out of every 10 satellites and those ones are ours and even though they share some infrastructure we decide what happens with those satellites", wouldn't surprise me at all if the Department of Defense was headed in that direction.

**What do you make of these LUCAS drones that are basically a rebuild of the Iranian ones but there's just there's like a Starlink on them or a Starshield on them or whatever it's called?**

**GA:** "Duh", would be my first reaction here. Ukraine and unfortunately, Russia, which was also accessing Starlink capabilities until they were, I think, more or less definitively shut off in recent months. But Ukraine has shown that, universal comms availability enabled by low Earth orbit satellites is really, really powerful and you can't always have that because there's jamming because there's electronic warfare because there's all these other things but in general satellite connectivity is like the lifeblood of the

force and so the United States — we have military satellite constellations, it's just that in general they're old and not really in the ultra-high bandwidth for a huge user population.

**And there's not that many of them so they could be shot down.**

**GA:** Yeah so like you literally get to the point where like the drone fleet over Afghanistan was like subscribing entire satellites for like one or two drones. And when I say satellites, I don't mean these like 100 kilogram Starlink LEO satellites, I mean like the 6,000 kilogram geostationary orbit satellites that cost $200 million a pop. So the bandwidth capabilities really had led the Department of Defense to a strategy of expensive, exquisite, even in the drone world and now what Starlink is enabling is this cheap, good enough, incredibly numerous type of capabilities. Because in the same way that we don't want $200 million satellites that are basically Fabergé Egg-level of fragile, and we don't want $4 million Tomahawk missiles, we want like $10,000 missiles.

**These LUCAS drones are $30,000.**

**GA:** Exactly. So this is this is all about the Department of Defense coming to terms with this this category of disruptive innovation, which is it's actually not as good

as what's in the arsenal right now. But it's so cheap that you can't ignore it and you can generate capabilities through mass that you might not have been able to generate with performance that you previously would have used performance to generate.

**And this is where I'm super sympathetic to, and the Trump administration ought to be more sympathetic to, the point about you really shouldn't just be going out trying to kill private companies, because this is obviously the way wars can and should be fought in the future. And it happened not because the U.S. government gave the idea of tens of thousands of low-Earth satellites that are possible because of reusable rockets, it's because a private company did and it's the same thing with AI.**

**GA:** Let me connect it to my own life story here, because I talked about the aftermath of Project Maven when Google pulled out in the summer of 2018.

**Google's just sitting on the sidelines, feeling pretty good right now.**

**GA:** And here's the thing. I spent so much of my life in the Department of Defense trying to convince Silicon Valley companies, "Hey, come on in, the water is fine, the defense contracting market, you know, you can have a good life here, just dip your toe

in the water". And what the Department of Defense has just said is, "Any company that dips their toe in the water, we reserve the right to grab their ankle, pull them all the way in at any time". And that is such a disincentive to even getting started in working with the DoD. And so, again, I'm sympathetic to the Department of Defense's position that they have to have control, but you do have to think about what is the relationship between the United States government, which is not that big of a customer when it comes to AI technology.

**That's the big thing. Does the U.S. government understand that?**

**GA:** No. Well, so you've got to remember, like, in the world of tanks, they're a big customer. But in the world of ground vehicles, they're not.

**To me, this is the real both sides talking past each other. I think that the U.S. has massively orders of magnitude less economic power than they think they do, and I think the tech companies dramatically underestimate the fact that they have guns and you don't, and that's why I wanted to write about this. Because what you're going to end up with is these differences are – the economic differences are only going to expand, and as the capabilities comes**

**up, the sense of loss in terms of having guns matters because there's actually a rival power base is going to – it's like the China and Taiwan thing. It's going to increase the sense we have to use the sticks that we still have, we didn't realize our carrots were baby carrots, and we don't want to get there, and it's not in anyone's interest to get there. And I just – again, it was not a fun week, I got called a fascist, called a bootlicker, all these sorts of things, but if it in some small way helps realize that all these concerns are real, but there's a much bigger problem here, which is a private company, like not literally at war, but not that far off, with the U.S. government. That's a much worse problem for everyone, and we have to avoid that.**

**GA:** Yeah, I certainly agree that the Department of Defense, to your point about like how if Anthropic ultimately becomes a rival basis of power with nation-state like significance, you can imagine why the DoD might want to get rid of them. I take your point.

But for me, the starting point has to be, Dario used to work at Google — he is, I think it's pretty obvious to everyone involved, probably sympathetic to the political left, right? And this guy has

converted this entire generation of folks who were skeptical about the national security establishment to enthusiastically supporting the military, what a gift he is in that regard. And so to demonize him is just so damaging to Silicon Valley national security relations in a way that I just cringe at. Because if you can get somebody like Dario enthusiastic about national security, if Dario can get thousands of Anthropic employees who, you know, five years ago were probably like in an EA group home, talking about pacifism and veganism and get those guys excited about national security — what a gift! What a gift that he has brought that community on board with literally developing autonomous weapons. Not using them, subject to no human oversight, but developing them, ready to go. This is so much progress, don't take us backwards.

**I will take that as a personal admonishment, I have not given Dario that degree of credit, and a note of optimism about a very interesting week. We had a whole section of the Iran War but we've already gone over just on this topic alone, but Gregory Allen is great to talk and I look forward to doing it in the future.**

**GA:** Yeah, thanks a lot.

This Daily Update Interview is also available as a podcast. To receive it in your podcast player, **visit Stratechery**.

The Daily Update is intended for a single recipient, but occasional forwarding is totally fine! If you would like to order multiple subscriptions for your team with a group discount (minimum 5), please contact me directly.

Thanks for being a supporter, and have a great day!

## Related

**2026.10: Higher Powers and Lower Macs**

Friday, March 6, 2026

**An Interview with Bill Bishop of Sinocism about COVID-19, U.S. China Relations, and Media Entrepreneurship**

Thursday, February 27, 2020

**An Interview with Dan Wang About China, the U.S., and Technology**

Thursday, April 1, 2021

**Anthropic's Skyrocketing Revenue, A Contract Compromise?, Nvidia Earnings**

**2026.10: Higher Powers and Lower Macs**

An Interview with Gregory Allen About Anthropic and the U.S. Government
Thursday, March 5, 2026

Anthropic's Skyrocketing Revenue, A Contract Compromise?, Nvidia Earnings
Wednesday, March 4, 2026

Technological Scale and Government Control, Paramount Outbids Netflix for Warner Bros.
Tuesday, March 3, 2026

View All

How High Can the Hornets Fly?, Book Club on Cavs-Pistons, Mail on Wemby, Desertion Rates, and Coach Flopping
Greatest Of All Talk | Mar 6


MacBook Neo
Dithering | Mar 6


(Preview) The Anthropic Mess Continues, Frontier AI and the Uncertain Future of Law, Q&A on Netflix, Dating Apps, F1
Sharptech | Mar 6


## Stratechery *Plus* INTERVIEWS

An Interview with Gregory Allen About Anthropic and the U.S. Government
Thursday, March 5, 2026

An Interview with Bill Gurley About Runnin' Down a Dream
Thursday, February 26, 2026

An Interview with Matthew Ball About Gaming and the Fight for Attention
Thursday, February 19, 2026

View All