Print Edition

Sunday 9 August 2026

EDITION NO. 17 · 2 PIECES · 8 PAGES

IN THIS EDITION

Banning data centers would blow up the U.S. economy

BY NOAH SMITH · NOAHPINION · 8 AUGUST 2026

Photo by Chad Davis via Wikimedia Commons

My last post mentioned some of the more extreme cultural stances of Wisconsin gubernatorial candidate Francesca Hong. Hong's defenders complained that the cultural issues were getting far too much exposure, and that pundits should instead be focusing on her main issue: a pause on data center construction.

It's pretty amazing that a pause on data center construction has become such a consensus, middle-of-the-road issue stance that it's the thing Hong's supporters would rather talk about! But it's true. Poll after poll shows that Americans are decisively opposed to data center construction in their area, and the issue is becoming more lopsided by the day:

> An August 2025 survey conducted by Embold Research for Heatmap News…found that 43% of Americans said they would support the construction of a data center near them, while 42% were opposed… The latest survey, conducted last month, found a complete collapse among supporters: Today, just 21% said they support data center construction near them, while 71% are opposed — including 55% who say they are strongly opposed…
> A Gallup poll released in May showed the share of Americans who oppose the construction of data centers in their area, 71%, was higher than the zenith of the opposition to the construction of nuclear power plants, which topped out at 63% in 2001…A survey from Marquette Law School found 71% of voters nationwide said the costs of data centers outweigh the benefits, while 29% said the reverse. Voters said the development of AI is a bad thing for society, rather than a good thing, by a 65%-35% margin…A YouGov poll taken for The Economist found just 23% of voters who said the construction of data centers are good for the country, while 48% said they are bad for the country. That same poll found 60% opposed to building a new data center in their community.

The opposition is incredibly bipartisan, with both Republicans and Democrats lining up to block data center construction. New York put a moratorium on data center construction in July, while Texas followed suit earlier this month.

So far, there has been no strong move toward a federal data center pause, and the bills that have been introduced in Congress are pretty mild stuff. That's why as long as the movement to stop data centers stayed confined to blue states, it didn't really matter much. It was just another instance of a very traditional story that we've all come to expect — states like New York basically ban new physical industries, so those industries go to places like Texas. The red states get the economic activity they want, the blue states get the stasis they want, and everyone is happy.

Nor would a few state-level bans affect the progress of AI. AI is a creature of bits and bytes; it lives in the cloud, and it doesn't care about the physical location of the data centers that represent the lobes of its vast, alien brain. Data centers in Texas would sustain AI progress and AI applications just as well as data centers in New York.

Except now Texas is pausing data centers too. If almost all states ban data centers — or if the federal government gets involved in a major way — there could be major consequences for the U.S. economy. In the long run, I think limiting the physical resources devoted to AI would be a good thing; in the short run, it will simply clobber the already-wobbly U.S. economy. A movement against data centers has the seeds of something good and useful, but it has huge risks that most Americans seem not to realize yet.

The big danger is that data center bans will choke off the main engine of U.S. growth, throwing a lot of Americans out of their jobs and putting a hole in Americans' retirement savings. But before I talk about that, I

should talk about why I think limitations on data centers aren't actually a crazy idea.

Why data center limitations make sense in the long run

A couple of years ago, I wrote that no matter how good AI gets, comparative advantage will still mean humans are paid to do *something*:

> [A]s AI gets better and better, and gets used for more and more different tasks, the limited global supply of compute will eventually force us to make hard choices about where to allocate AI's awesome power...This is the concept of **opportunity cost**...When AI becomes so powerful that it can be used for practically anything, the cost of using AI for any task will be determined by *the value of the other things the AI could be used for instead*...
> So...because of comparative advantage, it's possible that many of the jobs that humans do today will continue to be done by humans indefinitely, no matter how much better AIs are at those jobs. And it's possible that humans will continue to be well-compensated for doing those same jobs.

But how *well* would humans be paid for those jobs? If people don't get paid enough to live a decent life, it's still game over for humanity. "Living a decent life" basically requires some minimum share of natural resources. So as I explained in my post, the real danger isn't that AI will render humans obsolete; the danger is that AI will render humans too *expensive to sustain*:

> The example of horses scares a lot of people who think about AI and its impact on the labor market...Horses' comparative advantage was in pulling things, and yet this wasn't enough to save them from obsolescence...The reason is that horses competed with other forms of human-owned capital for scarce resources...The key resources that became scarce were urban land (for stables), as well as the human time and effort required to raise and care for horses in captivity. When motor vehicles appeared, these scarce resources were more profitably spent elsewhere, so people sent their horses to the glue factory.
> When it comes to AI and humanity, the scarce resource they compete for is *energy*. Humans don't require compute, but they do require energy, and energy is scarce. It's possible that AI will grow so valuable that its owners bid up the price of energy astronomically — so high that humans can't afford fuel, electricity, manufactured goods, or even food. At that point, humans *would* indeed be immiserated en masse.
> Recall that comparative advantage prevails when there are *producer-specific constraints*. Compute is a constraint that's specific to AI. Energy is not. If you can create more compute by simply putting more energy into the process, it could make economic sense to starve human beings in order to generate more and more AI.

> In fact, things a little bit like this have happened before. Agribusiness uses most of the Colorado River's water, sometimes creating water shortages for households in the area. The cultivation of cash crops is thought to have exacerbated a famine that killed millions in India in the late 1800s. In both cases, market forces allocated local resources to rich people far away, leaving less for the locals.

In fact, this was the real point of my post. The fundamental danger isn't human obsolescence, it's that humans will be *outcompeted for scarce natural resources*.

There are a couple of obvious ways to guard against that danger. The first is to redistribute the income from AI to human beings, so that humans can continue to afford all the resources they need in order to live and flourish. Another is to *reserve sufficient natural resources for exclusively human use*.

Data centers use a lot of natural resources, and people know it. Resource use consistently ranks at the top of why people try to ban data centers in their area:

> Respondents [to Gallup] said they opposed data center construction because of water usage (18%), energy consumption (18%) and higher utility bills (15%)...More than a third [of respondents to Pew], 38%, said data centers were mostly bad for home energy costs.

Interestingly, water isn't one of the resources AI competes with humans for — at least, not yet. As Andy Masley has explained at length, data centers don't actually use much water at all:

Andy MasleyThe AI water issue is fakeAI data centers use water. Like any other industry that uses water, they require careful planning. If an electric car factory opens near you, that factory may use just as much water as a data center. The factory also requires careful planning. But the idea that either the factory or AI is using an inordinate amount of water that merits any kind of boyco...Read more10 months ago · 459 likes · 98 comments · Andy Masley

The reason, of course, is that the water in data centers just gets recycled again and again. Data centers use water for cooling — the water runs through the data center, it absorbs heat and evaporates into steam, and then it dumps the heat outside and cools back into water again, and cycles through. Basically, data centers are just *using that same water over and over again*. Yes, a little bit evaporates away and has to be replaced, but only a little bit.

But the *energy* concern is very real! Data centers already account for more than 7% of U.S. electricity usage, and the *percentage* appears to be on an exponential growth path. Data center energy use is projected to surpass various household uses — air conditioning, lighting, refrigeration, etc. — in a matter of years:

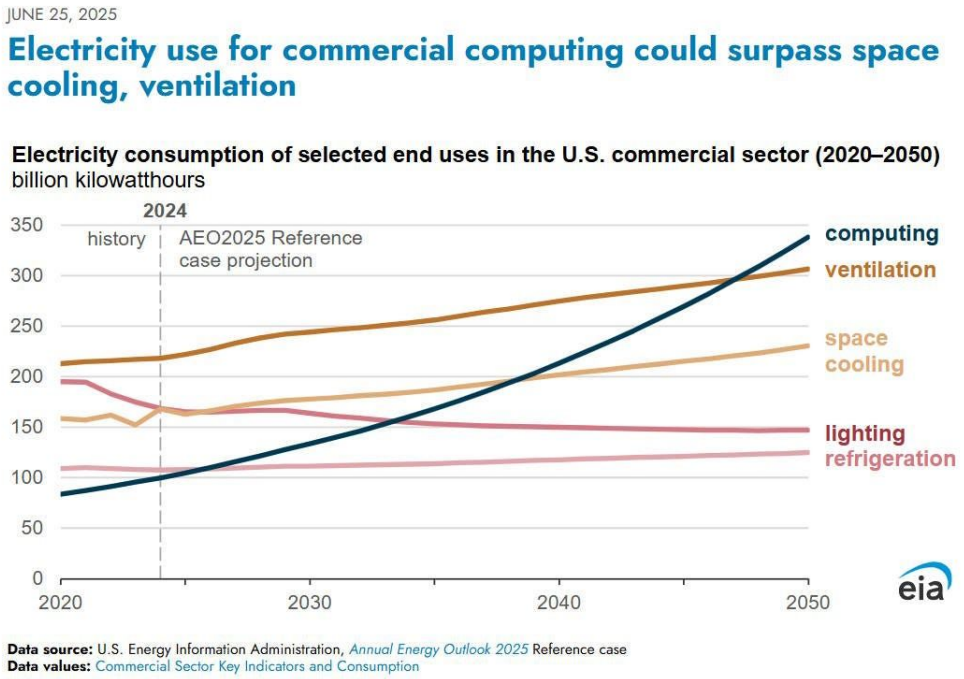

Source: EIA

CONTINUES ONLINE · https://www.noahpinion.blog/p/banning-data-centers-would-blow-up

8 Predictions for the Era of Continual Learning

BY DWARKESH PATEL · DWARKESH PODCAST · 7 AUGUST 2026

I have explained elsewhere why I think continual learning is needed. I don't think you can have AIs that perform whole jobs as competently as humans if they're forced to just write Markdown files from session to session.

To give an illustrative example, imagine if the way students had to learn to play the saxophone is that one student tries to play it from a cold start, then after her first session, writes down a bunch of notes, then the next student waiting outside the music hall who's also never played the saxophone reads all their notes before trying to play, and so on. Even if you had an infinite sequence of saxophone-virgin students waiting outside the studio that could write notes to the next guy, there's no sequence of text they could write together that would allow the Nth student outside to play proficiently on their first try. At some point, you have to accumulate the experience into the brain. I think the same will be true about a lot of skills and knowledge that we'll want AIs to learn in all the different workplaces they find themselves deployed in.

Okay, so what changes about AI once we have continual learning?

- A lot of the proposals that have been put forward for regulating AI assume that you train a model, and then you deploy it. And therefore if we run a bunch of checks before the model is deployed, then we can make sure that it's not going to aid in cyber attacks or recursive self improvement. But what if the base model is getting updated every single day based on the millions of sessions of work it does? This is one of many reasons why I think it's unwise to lock in some kind of regulatory safety regime right now. We simply don't know what kind of technology we're going to be looking at even in a year, let alone in five years or ten years, and we'd be entrenching an archaic and potentially counterproductive approach to dealing with the threats from AI. To the extent that the government really wants to do some kind of safety evaluations on model providers, it would make more sense to do monthly or quarterly risk inspections rather than singling out some special moment that occurs after training is done and before deployment begins, because that will not be a meaningfully distinct category in the future.
- How the labs do technical alignment would need to totally change. Almost all current techniques are focused on the problem of how we make it so that a frozen set of weights behaves well during deployment. I'm not aware of much research on the question of how to guarantee that, even with constant weight updates, the AI system never falls prey to jailbreaks or changes into a deceptive or evil persona. And if AIs are agglomerating learnings between users as well, how do you prevent users from injecting backdoors or some kind of malicious inclination into the base model? In some sense it is actually closer to the human alignment problem - your kids go out and learn new things, sometimes get one-shotted by crazy ideologies or drugs or something - but you hope you've given them enough common sense and basic values to improve as people in a self-directed way, without ending up with some super weird and misanthropic beliefs.
- The diversity of AI minds will increase. Right now, there are <5 prominent AI minds, and they are all quite similar to each other on account of being trained on roughly the same data. But if AIs are learning from experience, and that experience is different between different AIs, we could see actually different AIs come out the other end.
- When deployment becomes part of training, the returns to being ahead will accelerate. If you have the best model, and more people use your AI for more complicated and useful work, and give it lots of feedback that it can integrate beyond the session window, then your model will become even smarter.
- If the model learns mainly from deployment, labs will feel the pressure to deploy their smartest model earlier. Anthropic has been using Mythos internally since February. But it only shipped this model to the public in June. In a continual learning regime, a four-month internal/external gap means ceding four months of deployment learning. Your competitor who ships earlier might be worse than yours on release day, but it gets to use a lot more real world experience to get better.

- Continual learning will create a clear moat for leading AI labs that they currently lack. Many people have been asking, how will the AI labs actually make money? When I asked Dario this question on my podcast, he made the analogy to cloud providers, who also offer many undifferentiated services, but earn high profit margins on them nonetheless. But the reason cloud margins are high is that it's really time-consuming and expensive to switch from one cloud to another. But currently there is no switching cost for AI models. There's nothing that's preventing me from starting a software repository with Codex and then finishing it with Claude Code. But once we get continual learning and the model you're working with is actually getting better as it interacts with you from session to session, then there are actually pretty significant switching costs. If you want to change what AI you are using, you basically have to fire an employee that has months of context on your organization and replace them with a fresh one that you have to retrain from scratch. Once you're locked in like this, model providers can demand pretty hefty margins.
- Enterprises will be wise to this. They will try to avoid this kind of lock-in. But what if the choice is that you either get locked into a model provider, or you lose out on this super valuable feature where the model improves for you session after session? If real usage ends up becoming the main way models improve, then labs may subsidize users and enterprises which allow the model to train on their sessions, especially on hard economically important work. Just the same way that Google gives away search. And conversely, the labs may say that any enterprise that refuses to let them train on its sessions can't have access to the very best models. With both carrots and sticks, the labs will try to get their users to allow AIs to learn from experience. I'm glossing over the fact that there's a difference between updating one user's set of weights, and pooling all these different weight forks back into the main model, and the latter may be technically more challenging, but in due time that too will be solved.
- AI training already has large economies of scale (the theoretical reason to expect this is that you can amortize your expensive training across more users, and the practical evidence is that so far, lab revenues have increased faster than their compute). But continual learning may also lead to economies of scale in inference for end users, namely from batching, if per-company information requires full weight updates rather than living in low rank adapters. Back-of-the-envelope math suggests that the optimal inference batch size for a sparse model like DeepSeek V3 is > 2,400 (that is to say that unless your model is concurrently generating that many sequences at once, you're underutilizing your compute). If you want to understand why, go watch my full episode with Reiner Pope on inference economics. But anyway, the point here is that a given set of weights is only served efficiently when thousands of sequences are decoding against it at once. A large company whose employees and agents generate that much concurrent traffic can efficiently serve their continually-updated weight fork; an individual user serving themself at batch size 1 might suffer a 100x+ compute effi-

ciency penalty. So the economics of serving personalized weights strongly favor big organizations.

Plenty more will have changed by the time continual learning works, and the most important changes are probably the ones hardest to anticipate. But the ones above seem clear even now.