Google Data Center, Council Bluffs Iowa. Photo: chaddavis.photography
(CC BY 2.0)

If you'd asked me three years ago, I would never have guessed energy bills would be a salient political issue in 2025, let alone one that politicians anchor their entire campaign around. But of course, the cost of energy is quite different today than it was a few years ago. Electricity prices have increased faster than (already-high) inflation since 2022, leaving ratepayers on the hook for ever steeper monthly bills. At the same time, legislators are asking us to "electrify everything," or at least many things. Decarbonization cannot succeed at scale without increasing the demand on our energy grid—electric vehicles, heat pumps, and electric stoves are all necessary parts of a transition away from fossil fuels but are difficult to sell as electricity rates spike.

A confluence of factors has driven energy inflation but consumer advocates have recently pointed to data centers as a main culprit. The enormous energy demand of power-hungry AI training and inference centers, they claim, is forcibly raising costs for consumers. And although costs were already rising prior to Big Tech's capital expenditure bonanza, the influx of data centers seeking to connect to the grid has not helped. What's worse, the projected power demand for data centers is nowhere near its peak. Industry analysts estimate that U.S. data center demand will grow from 25 GW demand in 2024 to 80 GW by 2030. More than just additional generation is needed to sate that demand. New transmission and

distribution infrastructure will be required to connect data centers to the wider grid.

This week, [a coalition of more than 230 environmental groups demanded a national moratorium on new data center construction in the United States](), citing skyrocketing electricity bills as a main reason. That protest follows on the heel of [sixteen data centers whose construction had already been stalled as a result of local opposition]().

Ratepayers have reason to worry about impacts on their bills. First, ratepayers may be asked to partially pay for transmission, generation and distribution infrastructure that will do little to benefit them under traditional ratemaking structures. Second, the investments being planned today are meant to meet what many believe are overly optimistic projections of the AI industry's growth. If that growth fails to materialize, ratepayers may end up on the hook for billions in newly constructed assets that will never end up powering data centers.

Industry proponents argue that we should not worry too much about overinvestment in infrastructure. We have built excess infrastructure before and always end up finding use for it. [Analysts]() point to the railroad boom in the late nineteenth century, or the massive investment in fiber outlay during the dot-com bubble. Afterwards, both railroads and fiber contributed significantly to economic growth. The bubbles essentially directed private capital towards subsidizing public infrastructure buildout.

In the AI era, capital investment is bifurcated between rapidly depreciating assets and lasting assets. I would mostly group processing chips—GPUs and TPUs—in the former category. [While these are enormously expensive, they hold their value for only a brief period—essentially until the next generation of processing chips comes out](). Unlike railroads, which we were able to use for decades after their initial buildout, these chips are unlikely to be lasting infrastructure. On the other hand, we have physical infrastructure like transmission lines and power plants. This is the type of infrastructure that has potential to be useful outside of the AI context. Power lines care little about whether they are powering machine learning algorithms or an electric vehicle. And conveniently, [we will need a lot of infrastructure investment to meet the country's growing electricity needs regardless of whether the AI boom pans out](). Reshoring of industry and widespread electrification will structurally increase energy needs independent of data center

buildout. In an ideal world, then, the vast investments being made in new generation and transmission assets by tech firms will be productive regardless of whether they end up being used to power artificial intelligence.

However, this is true only insofar as infrastructure is built that can actually be repurposed for other uses. And right now, this is not always the case. For example, data centers are increasingly seeking to be co-located with generation facilities, meaning that a data center will be placed directly next to a generation asset. This allows data centers to sidestep the need for transmission or distribution investment, as they can hook directly into generation. Data centers can then argue that they do not really "use" the grid and try to avoid paying the cost of network infrastructure typically socialized across all ratepayers. What's worse, co-location offers few future benefits. No additional infrastructure is built, and reliable generation sources may enter into power purchase agreements with tech firms that take those generators away from the grid entirely.

A larger problem with the data center boom is that much of the infrastructure they are seeking to build is not the energy infrastructure we want. Gas turbines are faster to bring online, and offer attractive generation profiles for data centers, which require power 24/7. Gas turbines can also be built wherever and are easy to co-locate with data centers. The same can't be said for solar or wind, which may need to be built in a more remote area and be connected to load centers with high-voltage transmission links. Accordingly, we have seen requests by hyperscalers for utilities to construct more gas turbines. In Georgia, for example, a massive amount of new fossil fuel generation capacity is set to be brought online to meet data center load. In Louisiana, Meta entered into a specific agreement with Entergy to construct three new gas plants to meet its data center needs. Combined-cycle gas power plants can remain operational for up to 45 years. Once operational, gas plants are still cost-competitive with, if not cheaper than, solar and wind. Thus, we risk being stuck with an enormous amount of additional fossil fuel capacity for decades to come if regulators don't intervene.

What should intervention look like? For data center companies, time comes at a premium. Everyone is racing to deploy more infrastructure. It makes sense that waiting for the additional permits required to source clean energy (in the form of transmission and distribution infrastructure buildout) is a burden for these firms. The best way to deal with this remains streamlining permitting procedures for clean energy projects, battery storage, and transmission. Transmission is increasingly the bottleneck to clean energy deployment, so figuring out methods to

accelerate its construction will be key. Expedience is also part of the reason we see so many firms express interest in co-locating their data centers. But circumstances that promote co-location forego a chance to ask tech firms to subsidize grid buildout. The same is true for policies that provide an incentive for co-location, such as allowing co-located data centers to avoid grid infrastructure charges altogether. At the very least, regulators should require data centers to contribute to grid charges regardless of what form of interconnection they employ—as even utilities have asked regulators to do. We may also wish to discourage co-location with fossil fuel generation altogether. Placing data centers near renewable energy assets is a smart way of avoiding the need for transmission buildout—and has become a business model in and of itself—whereas moving data centers next to gas turbines provides us with few benefits.

Regulators should also recognize the leverage they hold over data center companies. Power is desperately needed, and regulators are in a place where they can impose conditions on firms seeking to buy it. Why not condition the permitting of a major new load center on the use of a certain percentage of clean energy? Further, regulators could require data centers to agree to flexible curtailment—one part of gas turbines' appeal is their consistent load profile, whereas renewables only produce at certain times of day. To alleviate this, model training loads—the energy used to train a model before it is released—could be batched and processed at times when grid load is lower or when renewable production is higher. While this may be less feasible for inference loads, model training accounts for between 10-20% of energy usage, and could easily be scheduled. This could reduce the need for costly peaker plants and smooth out the load curve for utilities.

The AI-driven capital expenditure boom is both an opportunity and a risk for clean energy advocates. On the one hand, we have a generational opportunity to push for rapid infrastructure development subsidized by some of the wealthiest corporations in the world and supported by an administration that is otherwise hostile to anything that reeks of renewable energy. On the other hand, we risk building out vast additional gas generation capacity that will remain active for decades and put us even farther off track to meet our decarbonization goals. Regulators should seize this moment to make sure we reap the benefits and, to the extent possible, avoid the costs.