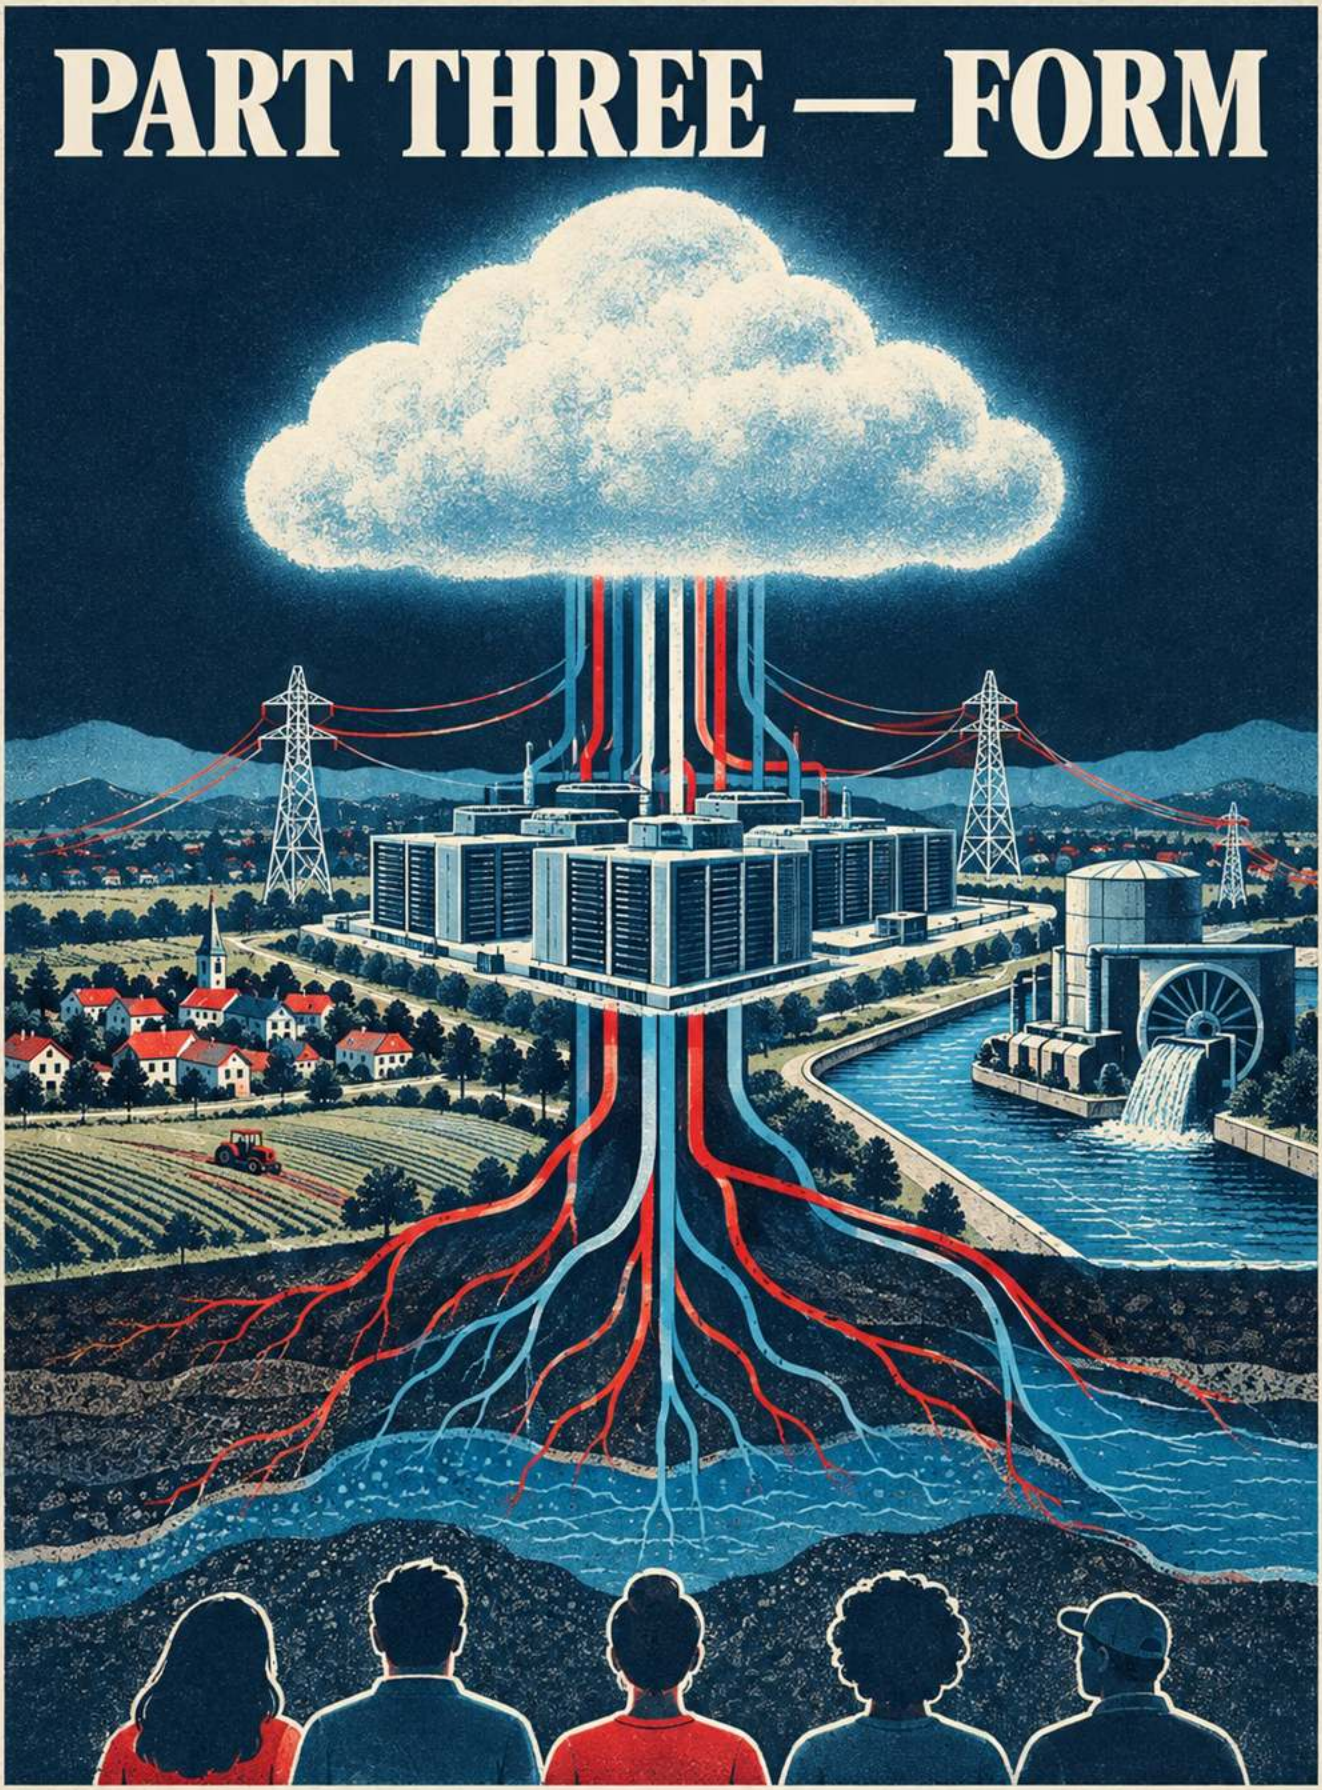
PART THREE — FORM



THE PHYSICAL BODY OF THE CLOUD

Energy, water, land, and the geography of who pays

A CITIZEN'S MANIFESTO FOR THE AGE OF AI

A Civic Blueprint for Human Agency, Democratic Self-Government, and a Workable Earth

PART THREE — FORM

THE PHYSICAL BODY OF THE CLOUD

Energy, Water, Land, and the Geography of Who Pays

DAVID M. DOYLE

The Series Charter

Book Zero—The Unoccupied Center—sets out the charter for this project: technology must enlarge human agency, not replace it. The six books that follow examine the conditions required to keep that commitment real.

Part Three concerns Form: the material integrity of the world beneath technological systems—land, water, air, climate, infrastructure, and community health.

The movement of the six-part manifesto is deliberate: See → Act → Protect → Share → Draw the Line → Build Together.

The Three Tests

Witness Test. Who sees this decision from outside the authority making it, and can their testimony survive?

Burden Test. Who bears the cost, and do they have meaningful voice, remedy, or refusal?

Return Test. How does this authority end, reverse, repair harm, or return value to those it affected?

How to Read Part Three

Artificial intelligence has a body: mines, servers, warehouses, data centers, power plants, transmission corridors, cooling systems, supply chains, tax agreements, workers, and places where heat, water demand, pollution, and risk are left behind.

Form asks whether full costs are visible, whether harms are prevented or repaired, and whether people who live with a project have real power to shape, condition, or refuse it.

Its standard is cumulative: no project should be judged only as an isolated permit when its actual effects accumulate across a watershed, grid, labor market, tax base, coastline, or community.

PART THREE — FORM

Structures that endure, and what they are permitted to hold

Its corruption: rigidity, control, and the treatment of a living place as a platform.

Contents

Chapter Seven — The Physical Body of the Cloud: Energy, water, land, and the geography of who pays

Chapter Eight — Climate and the Habitable Earth: Compute versus the carbon budget

Chapter Nine — Who Owns the Machine: Concentration, capital, and the four-company problem

Part Opening

Since 1932, coastal Louisiana has lost roughly 4,833 square kilometres of land — about a quarter of what it had — an area the United States Geological Survey compares, in its own report, to the land area of the State of Delaware.363 Three things happened to that coast in 2025.

On the twenty-fifth of April, the New Orleans District of the Army Corps of Engineers suspended the federal permit for the Mid-Barataria Sediment Diversion, the largest marsh-building project in the state's history — an engineered reconnection of the Mississippi River to the Barataria Basin, designed to let the river do again what the river used to do.364 On the seventeenth of July, the state terminated the project outright, reducing its authorized budget from $2.26 billion to the $618.52 million already spent, citing costs, permitting, and litigation, and describing the decision as resting on "a totality of the circumstances."365 And on the twentieth of November, the two federal agencies that decide which wetlands the Clean Water Act reaches proposed a definition under which, by their own written acknowledgment, "the proposed rule would define fewer waters and wetlands as within the scope of the Clean Water Act."366 So in the space of seven months the state stopped building land, and the federal government proposed to protect less of it.

I am not going to pretend the diversion was uncontroversial — the fishing communities that opposed it were defending a living, not an abstraction, and their objection was serious.

But the sequence is the subject of this part of the book.

This is where I live.

And it is where I learned that the most consequential decisions about a place are almost never made in a single room, by a single body, on a single record.

Let me say what I am not arguing, because Part Three is the part most easily mistaken for opposition to building things, and it is not that.

Form is not the enemy.

Form is what endures — foundations, grids, pipes, pads, levees, the rebuilt marsh itself.

A society that cannot build cannot shelter anyone.

The dam, the substation, the fiber trunk, the terminal, the ring levee: these are the material conditions of a decent life, and the instinct to protect a landscape by forbidding all intervention in it produces, in the end, a coast that erodes politely.

Building is the harder and more honourable act.

The discipline of form is narrower than a prohibition.

It is that a structure exists to shelter, never to possess.

A levee is legitimate because it protects the people behind it; it becomes something else when it is built to make one parcel developable at the cost of the parish downstream.

A data centre is legitimate infrastructure; it becomes something else when the aquifer beneath it is treated as an input and the parish around it as a site.

Part One asked whether you can see what is being done to you.

Part Two asked whether anyone can act on what you see.

Part Three asks a physical question: what is being permanently altered, who bears the alteration, and what will be left when the operator's balance sheet no longer requires the place.

The mechanism is not new to artificial intelligence.

What is new is the speed at which it is now being run.

Almost every environmental protection in American law is applied to a project.

Almost every harm that matters is produced by a region.

A single withdrawal, a single interconnection, a single fill of eleven acres of wetland is nearly always defensible on its own record, and the aggregate of forty such decisions in one watershed is very often not defensible at all — but there is no docket in which the aggregate appears, no hearing at which it is the subject, and after 2025, considerably less law requiring anyone to look.

Two changes, both in that year.

In February, the Council on Environmental Quality removed its own government-wide regulations implementing the National Environmental Policy Act — the rulebook that had told every federal agency how to scope an environmental review — effective the eleventh of April, 2025, leaving each agency free to use, amend, or narrow its own procedures.367 Then in May, the Supreme Court held that an environmental impact statement "need not address the effects of separate projects," that where an effect "arise[s] from a separate project — for example, a possible future project or one that is geographically distinct from the project at hand — NEPA does not require the agency to evaluate the effects of that separate project," and that an agency need not assess such effects "simply because those projects (and effects) might not materialize but for the project at hand, or are in some sense foreseeable." NEPA, the Court wrote, "is a purely procedural statute."368 I want to be fair to that reasoning, because as a description of what the statute says it is defensible, and the practice it corrected — litigation over reports as a substitute for litigation over decisions — was a real dysfunction.

But read it as a citizen standing in a parish rather than as a lawyer standing in a court.

The cumulative harm is now, as a matter of federal law, nobody's document.

There is no longer a place to ask the question.

Here is the pattern I would ask you to learn to recognize, because it will appear in all three chapters of this part.

A living place is converted into a platform.

Not by conquest and not usually by deception, but by a series of narrow, individually reasonable administrative findings, each of which treats the place as an input to something happening elsewhere.

Power is cheap here.

Water is available here.

The permitting timeline here is eight months rather than three years — and that last one is not a figure of speech: in December of 2025 the Louisiana Public Service Commission voted four to one to waive the requirements of its own general order for qualifying capacity additions and to target a Commission vote within eight months of application.369 Then set the scale against it.

American data centres consumed 176 terawatt-hours in 2023, about 4.4 percent of all electricity used in the country; the national laboratory that measures this projects a 2028 range of roughly 325 to 580 terawatt-hours, or 6.7 to 12.0 percent of forecast national consumption.370 Whatever else that is, it is the largest new industrial load of my lifetime, and it is arriving in the places least equipped to negotiate with it, at a moment when the federal requirement to consider it in aggregate has been narrowed twice in one year.

And the wetlands question underneath it has narrowed too.

In 2023 the Supreme Court held that the Clean Water Act reaches "only those wetlands with a continuous surface connection to bodies that are 'waters of the United States' in their own right," so that they are "indistinguishable" from those waters.371 The 2025 proposal would define that connection to require "surface water at least during the wet season and abutting (i.e., touching) a jurisdictional water" — the agencies acknowledging that this "might result in few wetlands being found to have a continuous surface connection under the proposed rule."366 A marsh that the state was, until July, spending public money to rebuild is a marsh whose legal protection is being measured, in the same year, by whether it touches.

Why does any of this belong in a book about artificial intelligence?

Because the cloud has a body, and the body is somewhere, and the somewhere has a water table and a school board and a hurricane season.

Parts One and Two concerned what a system does to a person's standing — what they can see, who decides, where they can go to argue.

Part Three concerns what the system does to the ground, which is the one harm in this book that cannot be appealed, because an aquifer that is drawn down or a marsh that is filled does not have a remedy, only a cost, and the cost is paid by whoever is still living there in thirty years.

That is also why this part carries the book's most conservative instinct.

I am willing to be argued out of a great deal.

I am not willing to accept a permanent alteration of a place in exchange for a reversible commercial benefit, on a record that was never required to add anything up.

So the chapters that follow ask four questions of every structure, and I would ask you to carry them into any hearing where a site is being discussed: What is being permanently altered.

Who pays if it fails.

What has been added up, and by whom.

And what is left when the operator leaves.

And here is the standard I would hold any large physical project to, stated as one test, because I have come to think the rest is engineering: a structure is legitimately sited to the extent that it is assessed together with everything else drawing on the same grid, watershed, and coast; that the money to undo it is posted before it operates rather than promised after; and that the place it occupies is left with no net loss of the commons it held before.

Each clause is falsifiable, which is why I state it that way.

A project reviewed alone, bonded never, and restored on a schedule enforceable only against a company that may not exist fails all three — and that description now fits a great deal of what is being built, quickly, on ground that public money is simultaneously trying to save. — D.M.D.

Chapter Seven: The Physical Body of the Cloud

Energy, Water, Land, and the Geography of Who Pays

Energy, Water, Land, and the Geography of Who Pays

When they say "the cloud," where is it, and what is it drinking? I. On the eighth of November, 2021, the city council of The Dalles, Oregon, adopted Resolution No. 21-028, approving an infrastructure agreement with two limited liability companies named Design, LLC and Moraine, LLC. The agreement runs for up to fifteen years. Under it, the companies would build water projects and convey to the city certificated groundwater rights "for a total of 3.88 mgd (6.01 cubic feet per second)," identified in the document by their certificate numbers — 46835, 46836, 46839, and a portion of 82480 — together with applications to the Oregon Water Resources Department "to convert use from industrial to municipal use."372 What is in it is a complete and honest picture of a small city's hydrology, stated by its own staff. The city's water system capacity was "9.9 million gallons per day (mgd) on a sustained average-day basis." The projected maximum-day demand at full build-out of the urban growth area, including the proposed development, was 17.5 mgd. And the aquifer: "While The Dalles is located within a state-designated critical groundwater area (meaning no new rights are available for use of native groundwater from The Dalles Groundwater Reservoir), the City has about 5 mgd of currently undeveloped groundwater rights from historical wells that have been abandoned, including the one in the basement of City Hall." The sustainable annual yield of that reservoir is given as 5,500 acre-feet, of which 3,700 acre-feet were then unused.373 A city inside a legally closed groundwater basin — no new rights available — was negotiating a large industrial water demand by having the industry bring rights with it, harvested from wells that had been abandoned, one of them under the city hall basement. That is not corruption. It is competent municipal engineering under a scarcity constraint, and I would rather have the engineers who wrote that memo than most of the people writing about this subject.

What is not in the document is the name of the company. Design and Moraine are the counterparties. The parent whose data centers those water rights would ultimately cool does not appear in the operative agreement. Years later, Oregon's own energy department would describe the same arrangement in plain terms — "Google invested $28 million to build a municipal water treatment and storage system," transferring the system and its water rights to the city and "providing the community with 100 million additional gallons of water capacity annually" — and would record that under a 2021 state Strategic Investment Program agreement the company committed to an annual payment equal to "the greater of 50% of what property taxes would have been or a minimum of $3 million per year," projected to deliver "over $45 million in direct payments" by the end of the abatement periods.374 The deal was, on its own terms, defensible: the company brought water rights and infrastructure and money, and the city got capacity it could not otherwise have bought. That is a great deal more than most host communities get. But hold the two documents side by side, five years apart, and notice what the public had to do to see the whole transaction. In 2021 a resident of The Dalles reading the council packet would have seen two LLCs, a water-rights conveyance, and a fifteen-year term. To learn who was on the other side of it, and what the tax arrangement was, and how much water the facility would actually consume, that resident would have had to wait for a state agency to write it down in a document that, when I fetched it, was still labeled a draft. That gap — between the physical thing being built and the record a citizen can read while it is being decided — is the subject of this chapter.

II.

Begin with the aggregate, because the aggregate is where the argument is usually won or lost, and because it has changed shape recently enough that most people's intuitions are calibrated to the wrong decade. For most of the 2010s the story of data centers was a genuine efficiency triumph. Lawrence Berkeley National Laboratory, in the most systematic federal accounting we have, records that United States data center electricity use "remained stable between 2014–2016 at about 60 TWh," and that by 2018 it had reached "about 76 TWh, representing 1.9% of total annual U.S. electricity consumption." Demand for computing exploded across those years; the energy did not, because workloads moved from inefficient closets into hyperscale facilities that were very good at their business.375

That era is over, and the report says when it ended and why: "In 2017, the overall server installed base started growing and Graphic Processing Unit (GPU)-accelerated servers for artificial intelligence (AI) became a significant enough portion of the data center server stock that total data center electricity use began to increase again." Those GPU-accelerated AI servers went "from less than 2 TWh in 2017 to more than 40 TWh in 2023." The compound growth rate went from "approximately 7% from 2014 to 2018, increasing to 18% between 2018 and 2023."376 By 2023, the total was "176 TWh," which the laboratory puts at "4.4% of total U.S. electricity consumption." For 2028 it publishes a scenario range rather than a forecast — "roughly 325 and 580 TWh in 2028," equal to "6.7% to 12.0% of total U.S. electricity consumption forecasted for 2028," which under a stated fifty-percent utilization assumption "would translate to a total power demand for data centers between 74 and 132 GW."377 That range is a scenario band, not a prediction. The low end and the high end differ by a factor approaching two, and the report says plainly that "the total range of energy use growth for 2024 and 2028 is highly dependent on the quantity and operation of AI servers." Anyone who tells you they know which end we land on is telling you about their priors. What is not a scenario is the break in the trend line. The Energy Information Administration records that between 2020 and 2025 United States electricity demand "grew about 1.7% annually compared with 0.1% annual growth between 2005 and 2019," and attributes it without hedging: "Electricity use by data centers is driving the electricity demand growth." Its regional numbers are the ones I would put in front of a state legislator, because they show this is not a national average problem — growth in ERCOT and PJM "averages 10% and 3%."378 In its own press release the agency is blunter still: "The driving factor behind this surge is increasing demand from large computing centers."379 Twenty years of flat demand ended. Utility planning, rate design, environmental review, and municipal water contracting were all built during the flat period, by people who reasonably assumed load growth was a solved problem. The institutions are running software written for a world that stopped existing around 2020. I

II.

Now the water, where almost every public argument I encounter — on both sides — is aimed at the wrong quantity.

The intuitive picture is a cooling tower evaporating a river. The direct number is real and rising: data centers consumed "21.2 billion liters of water" in 2014, and by 2023 hyperscale and colocation facilities accounted for "84% of the 66-billion-liter total," with hyperscale alone projected to consume "between 60 and 124 billion liters" by 2028. On-site water intensity, which had been "just over 0.36 L/kWh through 2023," is expected to rise to between "0.45 and 0.48 L/kWh" as chips get hotter.380 But here is the figure that reorganized my thinking, and I have not seen it in a single public hearing: "The total indirect water footprint of U.S. data centers

is nearly 800 billion liters, attributed to water consumed indirectly through electricity use, based on the regional electricity grid mix for U.S. data center locations." In 2023 the national average was "4.52 L/kWh of indirect water consumption," against a United States average across all electricity use of "4.35 L/kWh."381 Roughly sixty-six billion liters through the cooling system. Roughly eight hundred billion liters through the power plants. The water question and the energy question are the same question wearing different clothes, and a facility that proudly closes its cooling loop while drawing more megawatts has, in aggregate, increased its water footprint. A movement that wins on cooling water and loses on generation mix has won the visible fight and lost the material one. Then the second correction, which cuts against the alarmists — including, on my first draft, me. Virginia's Joint Legislative Audit and Review Commission examined the industry in the state that hosts more of it than anywhere on earth. Its finding: "In 2023, the data center industry used an estimated 2.1 billion gallons of water, with just over a third coming from reclaimed water instead of new withdrawals. Data center water use accounted for less than 0.5 percent of total state withdrawals."382 If your case rests on statewide volumes, you do not have a case in Virginia. Which is why the denominator is the whole argument. Two sentences later the same report records: "In 2023, 11 data center buildings each used over 50 million gallons, including one building that used 243 million gallons (10 percent of the industry's total use)."383 One building, a tenth of an entire industry's water in the largest data center market in the world. Water is not a national resource that happens to be distributed locally. It is a local resource that is sometimes summed nationally, and the sum is an artifact of the summing.41 Oregon's energy department makes the same point with the comparison a county commissioner can actually use. One Central Oregon facility "consumed about 70 million gallons in 2024," which the state puts next to "the 50 million gallons the Intel Aloha manufacturing campuses used in 2023," and equates to "about the amount of water one half of a center pivot in Central Oregon growing alfalfa consumed in the same year (about 70 irrigated acres), or about as much water as 420 Prineville households in 2021."384 Four hundred twenty households, or seventy acres of alfalfa. That is a real quantity, it is not apocalyptic, and whether it is acceptable depends entirely on the basin, the year, and who else is asking. Policy questions require a record, and the record is thin by design. Virginia requires a state permit only above "10,000 gallons per day from non-tidal surface waters, two million gallons per day from tidal surface waters, and 300,000 gallons per month from groundwaters," and because nearly all of these facilities buy from municipal systems rather than withdrawing directly, "only two data centers have their own DEQ withdrawal permits."385 The state does not know, in a permit file, what most of its data centers drink. It knows what its cities sold them.

IV.

Now the money, and I start by quoting the finding that most damages my own argument, because a chapter that hides it is not worth reading. The Virginia legislature's staff asked directly whether data centers are being subsidized by other ratepayers through rate design, and answered: "The study found that current rates appropriately allocate costs to the customers responsible for incurring them, including data center customers."386 That is the opposite of the claim you usually hear, it comes from a nonpartisan legislative audit body, and I accept it. The people who say the hyperscalers are not paying their electricity bills are, in Virginia at least, wrong on the record. The honest question was never whether they pay their bill. It is what happens to everyone else's bill because of the system that has to be built to serve them, and here the same report is unsparing: "Even though current rate structures appropriately allocate costs across customers, data centers' increased demand will likely increase system costs for all customers, including non-data center customers." It names three mechanisms: a large amount of new generation and transmission "that would not otherwise be

built, creating fixed costs that utilities would recover over the next several decades," a portion of which "would be paid by non-data center customers"; energy prices rising for everyone because supply cannot keep pace; and greater reliance on imported power, leaving the utility "more susceptible to spikes in energy market prices."387

The magnitude, for one utility's typical residential customer: "generation- and transmission-related costs increasing by an estimated $14 to $37 monthly in constant (or real) dollars by 2040 (independent of inflation)," with the report's scenario table putting +$37 at unconstrained demand and +$14 at half of unconstrained demand.388 Fourteen to thirty-seven dollars a month, forever, arriving as an ordinary rate increase with no data center named on the bill. That is the transfer. It is legal, it is the predictable result of correct cost-of-service ratemaking applied to a load that requires new fixed plant, and nobody has to intend it. There is a second, sharper exposure the same report identifies, and it is the one I would put in front of any state with partial retail competition. If large customers can leave the utility for competitive suppliers after the utility has built generation to serve them, "the fixed cost of any recently built generation is divided among the remaining customers." The utility's own estimate, if all currently eligible customers exercised that choice, is that "the cost-shift to other customers could exceed $600 million annually (a $150 per year cost impact for a typical residential customer)."389 The report's fifth recommendation asks the legislature to direct the utility "to develop a plan for addressing the risk of generation and transmission infrastructure costs being stranded with existing customers" and to file it with the commission.390 A recommendation to make a plan is a long way from a firewall. It is also more than most states have. And the wholesale market has already repriced. In PJM, the regional operator covering the Virginia corridor, the capacity auction for 2025/2026 cleared at "$269.92/MW-day," against a prior-year price of "$28.92 in RTO," with a total cost to load of "$14.7 billion."391 The following auction cleared at "$329.17/MW-day," with the Baltimore and Dominion zones at "$466.35/MW-day and $444.26/MW-day" — and the operator named the driver itself: forecasted peak load for the delivery year "increased year to year by more than 5,400 MW, driven largely by data center expansion."392 A tenfold increase in the price of capacity in a single cycle, and the grid operator naming the cause. The cost of keeping the lights on in that region has been repriced by this buildout, and it is collected from everyone in the zone.

V.

So what does it look like when somebody actually builds the firewall? Several states have, the terms are public, and this is the most useful thing in the chapter for anyone with a commission in their capital city.

In Ohio, the utilities commission approved a data center tariff for AEP Ohio on the ninth of July, 2025. The minutes are three lines long: "24-508-EL-ATA Ohio Power Company (Proposed Opinion and order) Approval moved and seconded, vote 5-0," with a commissioner recorded as describing the order as protecting ratepayers "from increased costs… and potential stranded cost."393 The terms as the utility proposed them, in its own sworn testimony, were these: data centers would pay minimum demand charges "based on 90% of their contract capacity under the proposed tariff, rather than 60% that would otherwise apply under the current general service tariff," with 95% for crypto and mobile data centers; contracts "for an initial period of not less than 10 years"; collateral "equal to 50% of the total minimum charges for the full term of the contract" unless the customer carried an A– and A3 credit rating and ten times the collateral requirement in liquidity; and an exit fee "equal to minimum charges for 36 months," available only after year five — which the witness spelled out as an effective floor of "eight years of 90% minimum demand charges." On a 500 megawatt customer, he calculated the minimum transmission charges alone at "approximately $40.2M annually."394 I have to flag a gap. I could not

retrieve the commission's actual order. The vote comes from the commission's own minutes; the as-filed terms come from the utility's testimony; and the as-approved terms — 85% minimum billing demand, a term equal to the ramp period "plus eight years," and the provision that "surety bonds are not an acceptable form of collateral" — I have only from the utility's own description of its own tariff.395 That last clause is worth pausing on: the utility refuses to accept the financial instrument that communities are told to accept when they ask for restoration guarantees. The people who understand this credit risk best will not take a surety bond for an electricity bill. Indiana produced the document I would hand to a commissioner who wants to see the whole architecture in one place. In Cause No. 46097, decided the nineteenth of February, 2025, the commission approved large-load terms applying to customers "with contract capacity greater than or equal to 70 MW at an individual site or 150 MW on an aggregated basis." The initial contract term is "not less than 12 years." Monthly demand "shall in no event be less than the greater of (a) 80 percent of the Large Load Customer's contract capacity… or (b) 80 percent of the… highest previously established Monthly Billing Demand during the past 11 months." Collateral is "equal to twenty-four (24) multiplied by" the maximum monthly non-fuel bill. Capacity reductions require "at least 42 months written notice." And the order requires the parties to reconvene on emergency load shedding "due to the Large Load Customers."396 Seventy megawatts, twelve years, eighty percent, twenty-four months of collateral, forty-two months' notice. A firewall with dimensions. DAVID M. DOYLE 144 And then the sentence that tells you what it does not do: "The Settling Parties agree that the Settlement Agreement does not limit the ability of the Commission or the Settling Parties to address cost allocation issues in a subsequent proceeding."397 The hardest question — who eats the stranded plant if the load does not materialize — was carried forward. Deferred, in the order that everyone will cite as having decided it. Wisconsin's commission in April of 2026 went the other way from the usual direction of regulatory travel: it "lowered the energy demand threshold for tariff eligibility from 500 MW to 100 MW," "extended the VLC tariff minimum initial term length to 15 years," and "removed a capacity-only option that would have allowed data centers to only pay for 75% of the costs."398 A commission tightening a utility's own proposal is rare enough to be worth naming. Texas legislated rather than adjudicated. Senate Bill 6 sets a demand threshold of "75 megawatts unless the commission determines that a lower threshold is necessary," requires "a flat study fee of at least $100,000" for initial transmission screening, and lets the commission demand "security provided on a dollar per megawatt basis."399 It also does something none of the tariffs do: it conscripts the buildout into grid reliability. After the grid operator "deploys all available market services," it "may direct the applicable electric utility… to require the large load customer to either deploy the customer's on-site backup generating facilities or curtail load," and such a deployment counts "as firm load shed" in reliability price calculations — with qualifying backup generation defined as generation "that, in the aggregate, can serve at least 50 percent of on-site demand."400 Read that as a citizen rather than an engineer. Texas has decided that the diesel and gas fleet sitting behind these buildings is part of the grid's emergency plan. That is either the most sensible provision in American large-load law or a statutory commitment to run hundreds of unpermitted-scale combustion units in the hours when the air is already worst, depending on how the standards are written. The bill does not say. The standards will, and they will be written in a proceeding almost nobody will attend.

VI.

Which brings me to the part of this chapter that made me angriest, and it is not about emissions or aquifers. It is about whether anyone outside the room is allowed to know the numbers. On the twenty-first of May, 2025, the Oregon Public Utility Commission decided, in Docket UE 430, whether the state's consumer advocate could

see the utility's data center load data. The information sought was not exotic: "CUST ID, BILLED KWH, BILLED KW, and FACILITY CAPACITY for PGE's data center load." The commission's disposition: "The confidential status of the information at issue is maintained."401

I want to give the commission its due, because its reasoning is more thoughtful than the outcome suggests. It said: "Having reviewed the pleadings filed, we find ourselves unpersuaded by both sides. We see the value of a more informed public policy conversation on these issues and want to foster such a conversation to the greatest extent possible. On the other hand, we are not ready to release the full suite of information that CUB requests, fearing that it will expose individual customer information." It did not dismiss the request; it directed the parties to confer on "how can customer size and usage data be characterized to support effective differentiation between loads and the risks and benefits they may bring to the grid without disclosing an individual customer's use information."402 That is a commission behaving in good faith. And the result is still that in the year the load is being built, the state's own advocate for households cannot see how much of it there is. The consequence is asymmetric in exactly the way that matters: the utility knows, the customer knows, and the party representing the people who will pay the $14 to $37 does not. Compare Louisiana, where the same question came out the other way. In the docket concerning Entergy Louisiana's agreement to serve a large data center in Richland Parish, an administrative law judge subpoenaed the company's investment and jobs figures and the amount of load the facility would need; the company moved to quash; and at the commission's session on the twelfth of August, 2026, counsel reported that "the Judge, on last Thursday, ruled that the Motion to Quash was denied." The matter was contested and carried over.403 I record it as a live proceeding rather than a precedent. But it establishes that the disclosure is obtainable. Texas, meanwhile, requires large loads to disclose their backup generation — "to the interconnecting electric utility or municipally owned utility," which must pass it to the grid operator.404 Disclosure to the institutions, not to the public. A defensible choice, and it should be named as the choice it is rather than described as transparency. And then the structural finding, from Virginia's legislative staff, that I think is the single most important sentence in this chapter for anyone who has ever sat through a rezoning hearing. Among its recommendations: that the legislature "expressly authorize local governments to (i) require proposed data center developments to submit water use estimates and (ii) consider water use when making rezoning and special use permit decisions related to data center development." Two more recommendations ask for express authority to require sound modeling studies and to set "maximum allowable sound levels for data center facilities," including low-frequency metrics.405

Understand what a recommendation like that concedes. In the largest data center market in the world, the legislature's own auditors judged it necessary to ask whether counties may lawfully require an applicant to say how much water it will use, or model how loud it will be, before approving it. The witness is not removed by conspiracy. It is absent by default, because nobody ever wrote the authority down, and the applicants are represented by people who have read the enabling statute more carefully than the board has.

VII.

The other thing these buildings do is burn fuel on site, and the permitting authority for that is often much smaller than the company.

In Memphis, the air permit for the gas turbines at a large AI facility was issued not by the Environmental Protection Agency and not by the state of Tennessee but by the Shelby County Health Department, which held a public hearing "April 25, 2025, from 5:00 p.m. to 7:00 p.m." on the application of "xAI CTC Property

LLC."406 The department later reported receiving "more than 1,700 public comments," said "each comment will be addressed before any decision is made regarding the permit application," and estimated it would "complete the permit review process in approximately sixty (60) days."407 Seventeen hundred comments to a county health department, on a two-hour evening notice, about combustion turbines for a computing facility. I could not locate the issued permit, or the department's engineering evaluation of it, on the department's own website. Advocacy organizations and party filings in the ensuing administrative appeal quote both by page and by number; I have chosen not to build the factual claims of this chapter on those filings, because a party's characterization of a document I cannot read is not evidence I am willing to hand you as fact. So the honest statement is this: seventeen hundred people commented on a permit whose text is not, as far as I can determine, posted for them to read. What is on the public record, from the municipal utility, is the scale of the grid connection at the same site: 150 megawatts energized, a second 150 megawatts requested, with the transmission upgrades and both substations built "at xAI expense," the first tranche split as "8MW of electricity from an adjacent, pre-existing MLGW substation" and "142MW… from a newly constructed electric substation." At a second site, the utility recorded that no final request had been made but that "a range of grid power options from 260MW to 1.1GW has been discussed."408

On water, the same utility document describes something genuinely constructive. A recycled wastewater plant, originally conceived to reduce aquifer draw at a power plant and a steel mill, got built faster because the company "offered to build a facility (at xAI expense) that uses discharged, treated wastewater from the Maxson Wastewater plant, to produce 13M gallons per day of recycled water to be used by xAI, TVA, Nucor Steel and others, significantly reducing consumption of water from the aquifer by those facilities," on a parcel the city sold it.409 And on electricity, "TVA/xAI/MLGW have a signed agreement requiring xAI to curtail their consumption of electricity from the grid when demand is high."410 That is what a good outcome looks like in this sector: thirteen million gallons a day off an aquifer, a curtailment obligation in writing, capital costs on the customer. It exists alongside a permit the public cannot easily read. Both facts are true, and any argument that needs one of them to be false is not an argument I want. I should also record the finding that cuts against the loudest version of the generator complaint. Virginia's auditors examined diesel backup generation across the largest concentration of it in the country and found it responsible for "less than 4 percent of regional emissions of nitrogen oxides and 0.1 percent or less of carbon monoxide."411 Less than four percent is not nothing in a nonattainment area, but it is not the story, and people who lead with it are spending credibility they will need later. What I would rather have is the boring instrument. The most useful text I found is a county zoning amendment in Pennsylvania that caps generator testing at "100 hours per year," requires that "all generators shall be Tier 4 diesel certified," and sets nighttime noise limits in both A- and C-weighted decibels, the C-weighted metric being the one that captures the low-frequency hum that actually reaches bedrooms. I have to label it honestly: the copy I obtained is a draft with its adoption date blank.412 An hours cap, an engine-tier requirement, and a low-frequency noise limit, in an ordinance, are worth more than any statement of environmental principle a company will ever publish.

VIII.

Then the land, which is the most permanent thing in this chapter and the least discussed. Federal wetlands permitting produces a public notice for each project, and those notices are the cleanest documents in this field.

Three from one Army Corps district, in the Virginia market: A five-building campus in Chesterfield County proposes to "impact 17.99 acres of palustrine forested wetlands (PFO), 8,115 linear feet of stream channel," with compensation of "43.31 wetland credits… and 9,213 stream credits," and avoidance of 121 of 147 acres of on-site wetlands.413 A project in Botetourt County on a 343.6-acre site proposes permanent impacts to 3.24 acres of wetlands and 6,715 linear feet of stream, mitigated by "approximately 5.09 wetland mitigation credits and 7,215 stream mitigation credits."414 A phase-two expansion in Louisa County proposes impacts to "4.20 acres of non-tidal wetlands and 6,302 linear feet of stream," mitigated by "7.99 wetland mitigation credits and 5,985 stream mitigation credits."415 Each is defensible. Wetland credits are a real market, mitigation banks are a real mechanism, and I would rather have compensated impact than uncompensated impact. Statewide, the same legislative report puts the footprint at "over 63 million square feet of data center" on "7,200 acres."416 The problem is not any one notice. It is that no document anywhere adds them up. Each is evaluated against its own site. And in May of 2025 the Supreme Court held that a federal environmental review "need not address the effects of separate projects," including a project that is "geographically distinct from the project at hand," and that an agency need not assess such effects "simply because those projects (and effects) might not materialize but for the project at hand, or are in some sense foreseeable."417 Here is what that means for a watershed: the seventeen acres and the four acres and the three acres are each assessed as if the others were hypothetical, and the aggregate a hydrologist would need to see is not any agency's responsibility to produce. Even Virginia's report notes, on the electricity side, that "both PJM and transmission owners must review the expected cumulative impact of demand growth on the transmission system, from proposed data centers and all other sources." Cumulative review exists in this sector. We do it for the transmission system, because the wires will physically fail otherwise. We do not do it for the wetlands or the aquifer, because those fail slowly and in somebody else's jurisdiction.

IX.

Now the ledger, because host communities are not being asked to accept these costs for nothing, and the benefits are real. Virginia's legislative auditors put the industry's economic contribution at "74,000 jobs, $5.5 billion in labor income, and $9.1 billion in Virginia GDP," including indirect and induced effects.418 Those are large numbers. They are also modeled numbers, and the direct-employment share of any data center is famously small relative to its capital cost, which is why I would treat that figure as a ceiling rather than a measurement. Against it, the price. Virginia's statutorily required biennial report on its data center sales-and-use tax exemption records exempt investment of $21,668,251,000 in fiscal 2024 with a tax benefit of $1,292,457,000 — and then, in fiscal 2025, exempt investment of $33,177,326,000 with a benefit of $1,941,390,000.419 A fifty percent increase in one year. The legislative auditors' framing is the one to remember: the industry "reported saving $928.6 million in sales taxes in FY23," of which the state portion was "an estimated $683 million, making it by far the state's largest economic development incentive, with the next closest incentive valued at $74 million."420 Fairness requires one more number, and it complicates my side of the ledger. The tax department's own estimate is that "eliminating the DCRSUT exemption would reduce net revenues by nearly $1.3 billion over five years."421 That is smaller than five years of the gross benefit, because it nets against economic activity the department assumes would not occur without the exemption. Whether that assumption holds is the entire question, and it is unfalsifiable in the direction that favors keeping the exemption — you cannot observe the facility that would have been built anyway. I think the assumption is probably too generous. I cannot prove it, and I am not going to pretend the number is fraudulent because I dislike where it points.

X.

And here is the hole in the floor. I asked, across state commission dockets, county ordinances, federal permits, and legislative reports, a simple question: who is required to post financial assurance for taking one of these facilities down, restoring the site, and removing what was built if the load never materializes or the business moves on? I found no state requirement. No public utility commission tariff among those examined here — Ohio, Indiana, Georgia, Wisconsin, Texas — requires a decommissioning bond. Every financial assurance provision I quoted in section five secures unpaid electricity bills, which is the utility's exposure, not the community's. The Army Corps notices impose mitigation credits for wetlands, which is compensation for the impact of building, not assurance against abandonment. The Virginia report does not raise it. The only instruments I located were local, and both were unenacted or pending when I read them. A Colorado county ordinance provides that "the decommissioning obligation will be secured by a letter of credit, bond or cash deposit, in an amount based on a certified estimate prepared by a professional engineer," revised every five years.422 The Pennsylvania draft requires decommissioning to begin within a year and finish within eighteen months.423 That is the entire state of the art: two county documents, at least one a draft. DAVID M. DOYLE 150 Set that beside the interconnection queue. The Texas grid operator told a legislative committee that it "is tracking approximately 410 GW of Large Loads seeking interconnection of which ~87% are data centers," alongside "2,008 active generation interconnection requests totaling 453,562 MW."424 Queue volume is a pipeline of requests, not committed load, and the operator says so; most of it will never be built. That is the point. An industry whose own pipeline is understood by its own grid operator to be several times larger than what will actually exist is an industry with a very high expected rate of abandoned projects, and there is no bond behind almost any of them. We have built, in twenty-four months, a regulatory apparatus that carefully secures the electricity bill and does not secure the building.

XI.

The concerns, in the form I would bring to a hearing. The left column is the worry; the middle is the document to ask for, by name, on the record; the right is who actually decides it, which is usually not the body people show up to. If you get honest answers to the middle column, you do not need me. If you cannot, that is the finding, and it is reportable on its own.

XII.

The three tests, applied to this concentration. Witness. A commission that maintains confidentiality over billed kilowatt-hours and facility capacity, against its own consumer advocate, has removed the witness — however carefully, however provisionally, however much it invites the parties to confer.401 A county that is not clearly authorized to require a water-use estimate before it votes has never had a witness to remove.405 A permit that draws seventeen hundred comments and is not posted where those seventeen hundred people can read it has converted witness into ritual.407 The test is not whether officials are sincere. It is whether a specific document, containing a specific number, is available to a specific person before the decision is final. Burden. Follow the fourteen to thirty-seven dollars.388 It is the cleanest illustration in this book of a burden that is nobody's fault: the rates correctly allocate costs, the utility correctly builds to serve load, the commission correctly approves prudent investment, and a household in the same zone pays more every month for forty years with no mechanism by which its consent was ever sought. If the load fails to materialize, the fixed plant remains, and the

tariffs I quoted secure the electricity bill, not the plant.397 If the site is abandoned, no bond exists.422 The asymmetry is not in anyone's intentions. It is in the instruments. Return. Almost nothing here is designed to be given back. Wetland credits compensate for a conversion; they do not restore the parcel. A tax exemption is permanent by design. Water rights transferred from a closed basin do not return to it. The single genuine return I found is the Memphis recycled-water plant, which takes thirteen million gallons a day off an aquifer that a city and a power plant and a steel mill all depend on.409 It is real, it was voluntary, and it was faster than the public process would have been — which is exactly why it cannot be the model. A return that depends on a company's willingness is a gift, and gifts are revocable.

XIII.

What I would insist on. Six conditions, each with a name attached and a consequence for failure.

1. Published load and water, per site, before approval and annually thereafter. Contract capacity in megawatts, the ramp schedule, maximum-day and maximum-hour water withdrawal by source, and the identity of the customer taking service. Not aggregated to the utility level, not summarized to a percentage, not available only to parties under protective order. If a facility's load is a trade secret, then its cost is a public subsidy of a trade secret, and the state should decline to build the transmission. I would accept one narrow redaction — the internal allocation of compute among a customer's own business lines — and nothing else. Who: the state utility commission, by order in the tariff docket; the municipal water utility, by term in the service contract; the county, as a condition of the special use permit. If not: no certificate, no service agreement, and no permit issues; and any agreement executed without the disclosure is voidable on petition by any ratepayer or adjoining landowner.

2. A ratepayer firewall with dimensions, not a principle. The Indiana order is the template because it is specific: a megawatt threshold that triggers the class, a minimum contract term of at least twelve years, a minimum billing demand at a stated percentage, collateral measured in months of billing, and a notice period for reductions measured in years.396 To that I would add the term every one of these instruments is missing: an allocation, decided in the same proceeding, of who bears the cost of generation and transmission built for a load that does not arrive. Not deferred to "a subsequent proceeding."397 Decided in the order that approves the load, because that is the only moment at which the applicant has any reason to accept it. Who: the utility commission; the state consumer advocate, funded and appearing as a party of right; the legislature where the commission lacks authority. If not: the tariff does not take effect, and cost recovery for the associated plant is presumptively denied in the next rate case.

3. Peak-withdrawal caps with drought and heat triggers, written as numbers. Annual averages are the wrong instrument for a resource that fails on the worst day of the year. Every large-load water contract should state a maximum daily withdrawal, a maximum hourly rate, the source of each increment, and a mandatory curtailment schedule keyed to a published drought index or reservoir elevation — with the facility's reduction stated in gallons, not in "commercially reasonable efforts." Texas has already established that these buildings can be required to curtail electricity in a grid emergency.400 There is no principled reason water emergencies should be treated as less real. Who: the municipal or district water utility; the state water resources agency issuing the withdrawal permit; the basin authority where one exists. If not: service is reduced to the contract minimum by operation of the contract, without a hearing, and the facility bears the cost of any alternative supply.

4. Cumulative review across all pending projects in a region, restored by statute where the courts have removed it. After the 2025 holding that a federal review need not address geographically distinct projects, cumulative assessment will not come from federal environmental law.417 It has to be legislated at the state and regional level, and the model already exists on the electric side, where the grid operator and transmission owners are required to review "the expected cumulative impact of demand growth… from proposed data centers and all other sources." Do for the watershed and the airshed what we already do for the wires: one assessment, all pending applications, published before any of them is decided. Who: the state environmental agency; the regional planning district or basin commission; the legislature that must grant the authority. If not: individual permits are stayed until the regional assessment is published.

5. A decommissioning bond, sized by an independent engineer, posted before the first foundation. This is the requirement with no incumbent to defend it, because at present almost nobody has it. The Colorado county language is the right shape: security by "a letter of credit, bond or cash deposit, in an amount based on a certified estimate prepared by a professional engineer," recalculated every five years.422 The estimate should cover demolition, removal, restoration of grade and hydrology, and the cost of the utility interconnection nobody else will want. And the instrument should be one the beneficiary would actually accept — noting that the utility in Ohio, protecting its own revenue, will not take a surety bond.395 Communities are entitled to the same standard of care a utility applies to its own money. Who: the county or municipality issuing the land use approval; the state environmental agency as co-beneficiary; an engineer selected by the locality and paid by the applicant. If not: the certificate of occupancy does not issue, and operation without a current bond is an ordinance violation accruing daily.

6. Enforceable community benefit, with the abatement conditioned on it. Where a facility receives a tax exemption or abatement, the value of that exemption should be stated in dollars in the approving document, the term in years, and the benefits offered in exchange — payments, water infrastructure, workforce commitments, noise limits — written as covenants running with the land, with the abatement automatically reduced pro rata for each unmet commitment. A state whose largest economic development incentive is nine times its next largest should be able to say, per facility, what it bought.420 Who: the local governing body; the state revenue department publishing per-recipient values; the economic development authority administering the agreement. If not: the exemption terminates prospectively and the abated amount for the year of the breach becomes due. If you can carry only two, carry the first and the fifth. The first, because every other requirement in this chapter is unenforceable without it. You cannot cap a withdrawal you are not told about, or firewall a cost you cannot see, or assess a cumulative impact whose components are confidential. And the fifth, because it is the only one that changes the arithmetic before the decision rather than after it. A bond makes the company's projected build cost include the possibility of its own failure. Nothing else on this list does that, and no company in this sector has ever had to price it.

XIV.

I have to end with a confession, because I was on the other side of this table and it would be dishonest to write a chapter about disclosure without saying what I did about it. I bought capacity. Not at hyperscale — my company was a customer, not a builder — but I signed the agreements, and I remember exactly how the confidentiality worked, because I asked for it. We did not want competitors to know our capacity commitments, because capacity commitments imply growth rates, and growth rates were the number our investors cared about

most. So we required that our loads be treated as confidential, and our vendors were happy to agree, and every one of us in that negotiation would have described the arrangement as ordinary commercial hygiene. Nobody in the room was hiding anything from a county. And yet the aggregate effect of thousands of us behaving that ordinarily is the situation I described in section six: a state utility commission that cannot release billed demand to its own consumer advocate without fearing it will "expose individual customer information," where the individual customer is a corporation and the information is how much of a shared public system it has reserved. I also remember how we chose locations, and it was not complicated. Cheap power, tax treatment, fiber, water, and — though nobody phrased it this way in a meeting — the absence of friction. Friction meant a county with a professional planning staff, an active neighborhood association, a hydrologist on retainer, a history of litigating permits. We did not call those places hostile. We called them slow, and slow was a cost, and the site selection model priced it. A community's capacity to ask questions was, in my industry's actual practice, a locational disadvantage. That is the thing I most want a reader to take out of this chapter, because it survives every good-faith actor being good-faith. The buildout is not seeking out places with weak institutions in order to harm them. It is seeking out places where the process is fast, and process is fast where institutions are thin, and so the map of this industry is, with real exceptions, a map of where the fewest questions get asked. The remedy is not to be against the buildings. I am not. The computation is going somewhere, the demand is real, and a country that refuses to host any of it will simply export the emissions and the aquifer draw and the jobs to somewhere with fewer protections and then congratulate itself. Digital infrastructure is legitimate, and in some form necessary. What it is not entitled to is subsidized power, unmetered water, waived review, or the quiet transfer of its risks onto the people nearest to it. Those four are not incidents of the industry. They are, at present, its business model in the places where nobody made it otherwise. And the thing that makes it otherwise is embarrassingly small. Not a movement. A number on a piece of paper, in a file a resident can open, before the vote. The city staff in The Dalles wrote down the aquifer's sustainable yield, the unused portion, the system's daily capacity, and the build-out demand, in a public packet, before their council voted. Whatever else you think about that transaction, that is what it looks like when a community negotiates from knowledge instead of hope. They knew what they had, so they knew what to ask for, and they got water rights and a treatment plant and a payment floor out of it. Every other community in this chapter is being asked to do the same deal without the memo. — D.M.D.

Chapter Eight

Chapter Eight: Climate and the Habitable Earth

Compute versus the carbon budget

Are we spending the remaining carbon budget on the tool we claim will solve the carbon problem?

I.

There is a paragraph in the Code of Federal Regulations that settles, for one narrow purpose, the argument this chapter is about. It appears in the Treasury regulation governing the clean hydrogen production credit, and it defines a term. "Temporal matching," it says, "means that the electricity used by the hydrogen production facility is produced by the electricity generating facility in the same hour that the electricity is used by the hydrogen production facility." Then it grants a delay. "For taxable years beginning before January 1, 2030, a taxpayer may demonstrate temporal matching on an annual basis. For taxable years beginning after December 31, 2029, a taxpayer must demonstrate temporal matching on an hourly basis."425 The United States government, writing a rule that would decide who received billions of dollars, considered whether a company could claim to be running on clean electricity by buying a year's worth of clean generation somewhere on the grid and setting it against a year's worth of consumption. It decided that it could not — that the electricity has to be generated in the same hour it is used. Because the industry said it was not ready, the government wrote annual matching into the rule as a transition, with an end date. Treasury described the change in its own announcement of the final rules: the transition allowing annual matching was extended "two additional years relative to the proposed rules, with hourly matching required starting in 2030 for all facilities."426 So the position of the federal government, in binding regulation, is that annual matching is a temporary accommodation and hourly matching is the truth. For hydrogen.

For everything else — for every data center, every corporate net-zero pledge, every press release announcing that a company runs on one hundred percent renewable energy — annual matching remains the standard, because the body that writes the accounting rules has not finished replacing it. The Greenhouse Gas Protocol's technical working group has proposed exactly the reform Treasury already imposed: "requiring organizations using certificates to match them to consumption hourly, and from deliverable grid regions." The proposal cleared the Independent Standards Board for consultation. And then the timetable: "Following approval through the GHG Protocol governance process and publication of the revised Scope 2 Standard (anticipated late 2027), implementation is expected to phase in over multiple years."427 The consultation opened on the twentieth of October, 2025, was scheduled to close on the nineteenth of December, and was then extended: "The public consultation period for both proposed updates has been extended to January 31, 2026 in response to stakeholder feedback." A second consultation would follow in 2026, "with final publication of the new standard expected in 2027." Among the concessions already built into the package: "the use of load profiles to approximate hourly data, exemption thresholds for which organizations are covered, a legacy clause for existing contractual commitments, and a multi-year phased implementation timeline."428 Every megawatt of the buildout described in the last chapter is being contracted, permitted, financed, and poured now — under an accounting rule that its own standard-setter is in the middle of replacing, with a legacy clause for the contracts signed before the replacement lands. The concrete will be scored under the old rule. The claims made about it

will be, too. By the time the honest rule takes effect — late 2027 at the earliest, phased in over years after that — the plants will be built, the twenty-year agreements will be signed, and the carbon will be in the air. That is not a scandal. Nobody hid it. It is written down in the standard-setter's own bulletin and in the Federal Register. It is how a fast industry and a slow rulebook interact, and it is the single most important fact in this chapter.

II.

Now the quantity that the rulebook is trying to measure.

The Intergovernmental Panel on Climate Change put it in one sentence in its synthesis report: "The best estimates of the remaining carbon budgets from the beginning of 2020 are 500 GtCO2 for a 50% likelihood of limiting global warming to 1.5°C and 1150 GtCO2 for a 67% likelihood of limiting warming to 2°C." The same document sets the historical total against it: "Historical cumulative net CO2 emissions from 1850 to 2019 were 2400 ± 240 GtCO2 of which more than half (58%) occurred between 1850 and 1989, and about 42% occurred between 1990 and 2019." And it includes a finding that ought to end a certain kind of argument: "Projected CO2 emissions from existing fossil fuel infrastructure without additional abatement would exceed the remaining carbon budget for 1.5°C (50%)."429 Existing infrastructure. Not new. The plants already standing, run out their expected lives, overshoot the budget by themselves. Against that, the annual flow. The International Energy Agency reports that global energy-related carbon dioxide emissions rose about 0.4 percent in 2025 — "the slowest rate since 2021" — and that "total energy-related CO2 emissions increased by around 145 million tonnes (Mt) in 2025, reaching a new high of nearly 38.4 billion tonnes (Gt), and 5% above 2019 levels."430 Do the division yourself. Five hundred billion tonnes, from the start of 2020, at something under thirty-nine billion tonnes a year from energy alone. Thirteen years, and six of them already spent. That arithmetic is mine, not the panel's; it ignores non-energy emissions, it ignores the possibility that the rate falls, and it treats a probabilistic budget as a fuel gauge. None of those objections changes the order of magnitude. The budget for the more ambitious target is roughly a decade of current behavior, and we are most of the way through it. Which brings me to the number that ought to be the headline of the American energy conversation and is not. The Energy Information Administration reports that "U.S. energy-related CO2 emissions increased by 2%, or about 115 million metric tons (MMmt), in 2025," to 4,904 MMmt. And within that: "CO2 emissions from the electric power sector rose by 4%, or 58 MMmt, in 2025, partly because of growth in overall electricity generation." The agency names the causes: "Hot summer weather, which led to increased demand for space cooling and record summer peak electricity demand, and growth in electricity demand from data centers and manufacturing facilities both contributed to increased generation."431 The power sector was the part of the American economy that was reliably decarbonizing. In 2025 it went the other way, and the agency that keeps the books lists data centers among the reasons in the same breath as a hot summer.

The generation total tells the same story from the other side: "In 2025, the United States generated 4.43 thousand terawatthours (TWh) of electricity, up 2.8% from 2024 generation, which previously had been the highest annual total in our Monthly Energy Review Data Browser dataset dating back to 1949." The driver, in the agency's words: "Much of the recent growth in electricity demand comes from the commercial sector, which includes data centers, and the industrial sector, which includes manufacturing establishments."432 And the projection, from the national laboratory report I relied on in the previous chapter: data centers reached "176 TWh by 2023, representing 4.4% of total U.S. electricity consumption," with 2028 estimates spanning "roughly 325 and 580 TWh," or "6.7% to 12.0% of total U.S. electricity consumption forecasted for 2028."433 A spread

from 6.7 to 12 percent of national electricity is an admission that nobody knows. But note which direction the uncertainty runs. The low case is a near-doubling.

III.

Here is where the accounting question and the physical question meet, and here I have to give credit to two companies that could have said nothing.

Google reports that in 2025 it navigated "our largest load growth in history — a 37% annual increase in electricity demand," reduced operational emissions by two percent, and matched "100% of our electricity consumption with renewable energy purchases for the ninth consecutive year." Then, in the same document, it says the thing the press release does not usually say: "Our AI infrastructure buildout is accelerating faster than the grid is decarbonizing, and long waits to connect to the grid, fragmented markets, supply chain delays, and regulatory bottlenecks continue to slow down new carbon-free energy from coming online." And: "Although we were successful in decreasing our operational emissions, our supply chain emissions grew by 25% year-over-year."434

That is a self-report against interest, and the company had already made the analytical point in an earlier report with unusual clarity. "Achieving 24/7 CFE is a far more complex and technically challenging pursuit than annually matching our energy use with renewable energy purchases," it wrote — and then: "despite achieving a 100% global renewable energy match in 2023 for the seventh consecutive year, our reported Scope 2 emissions increased." The company explains why it is moving anyway: "The even more granular local and hourly accounting of 24/7 CFE provides a more accurate representation of corporate carbon footprints, and research has found that targeting 24/7 CFE drives greater impact on grid-level decarbonization than other procurement goals, such as a 100% annual renewable energy match." Its goal is to "run on 24/7 carbon-free energy on every grid where we operate by 2030," and its stated position in 2023 was "a global average of approximately 64% carbon-free energy across our data centers and offices."435 Sixty-four percent, on the honest measure, in the same year as one hundred percent on the official one. The following year, by the company's own account, "our carbon-free energy use from 64% to 66% on an hourly basis."436 Now Microsoft, whose disclosure is the most useful document I read for this chapter. Set its base year against its most recent year, from its own data sheet. Scope 1: 118,100 tonnes of carbon dioxide equivalent in fiscal 2020, 170,887 in fiscal 2025. Scope 3: 12,487,000 in fiscal 2020, 18,243,000 in fiscal 2025. And Scope 2, market-based: 456,119 in fiscal 2020 — and 2,707,428 in fiscal 2025.437 A market-based Scope 2 number that rises by a factor of roughly six. The company's own explanation: total emissions "increased 25% year over year, driven primarily by the expansion of our datacenter infrastructure and pausing our use of non-additional, unbundled renewable energy certificates as we prioritize investments that bring net new power to grids." And: "Scope 3 remains the largest share of our footprint overall, but one of the clearest changes this year was the growing contribution of Scope 2, which represents 13% of our total emissions—up from nearly 2% last year."438 Part of the increase is not new pollution. It is the removal of an instrument. The company stopped buying certificates it describes as "non-additional" and "unbundled" — certificates that did not cause any new clean generation to exist — and its reported emissions went up by millions of tonnes. The earlier, lower number was partly an artifact of the instrument, not a description of the electricity. The atmosphere did not change when Microsoft changed its purchasing policy. Only the ledger did.

I hold this up as an example of honesty, because a citizen's manifesto that cannot distinguish between a company disclosing an uncomfortable number and a company hiding one is useless. Microsoft's long-stated ambition is stricter than the standard it is scored against: "By 2030 Microsoft will have 100 percent of our electricity consumption, 100 percent of the time, matched by zero carbon energy purchases."439 One hundred percent of the time is the hourly test. The company set it for itself six years before the standard-setter proposed it. What follows is not that these firms are lying. It is that the public number and the physical number are different numbers, that the gap between them is large, and that until 2027 at the earliest the public number is the one the rulebook blesses.

IV.

So look at the rulebook itself, since so much rests on it. The 2015 Scope 2 Guidance requires dual reporting: "Companies with any operations in markets providing product or supplier-specific data in the form of contractual instruments shall report scope 2 emissions in two ways and label each result according to the method: one based on the location-based method, and one based on the market-based method."

The location-based method is the physical grid. The market-based method is the certificate portfolio. Both are supposed to be published. The operative rules for the market-based method are five quality criteria. Certificates must convey the emission-rate attribute, be the only instruments carrying that claim, be retired on behalf of the reporting entity, "be issued and redeemed as close as possible to the period of energy consumption to which the instrument is applied," and "be sourced from the same market in which the reporting entity's electricity-consuming operations are located."440 "As close as possible." "The same market." Those are the two criteria that hourly matching and deliverability would replace, and in the surrounding guidance the language softens: instruments "should be sourced from regions reasonably linked to the reporting entity's electricity consumption." Should, not shall. And then the standard says the quiet part about itself. On residual mix: "If a residual mix is not currently available, reporters shall note that an adjusted emissions factor is not available or has not been estimated to account for voluntary purchases and this may result in double counting between electricity consumers."440 The accounting standard concedes, in its own text, that the same clean electricity may be counted by more than one buyer. That is a disclosure requirement.

The Environmental Protection Agency describes the same gap in plainer English: "Consumers have most often purchased electricity generation to match the consumer's annual electricity consumption on a total annual volume basis," and "while the volumetric annual matching of generation plays an important role in helping scale demand for new zero emissions resources, it may not fully realize a power sector that provides electricity from clean energy all the time." The agency describes hourly matching as an approach that "typically involves buying the electricity from the same regional area where the consumer's electricity consumption occurs" and "seeks to optimize the impact of procurement from a time and space perspective."441 Meanwhile the consumer-protection rules are narrower than most people assume. The Federal Trade Commission's Green Guides address renewable-energy claims directly. It is deceptive to make an unqualified renewable energy claim "if the claim fails to disclose that the renewable energy was not used to manufacture any part of the advertised product or package." A marketer "should not make an unqualified renewable energy claim, directly or by implication, if the claim is based on the purchase of renewable energy certificates (RECs) without also disclosing that the marketer purchased RECs. For example, if a marketer generates renewable energy, but sells RECs for all of the renewable energy it generates, it would be deceptive for the marketer to represent, directly or by implication,

that it uses renewable energy."442 So selling the certificate while claiming the electricity is addressed. Matching a July evening in Louisiana against a windy March night in Oklahoma is not. And the Guides are weaker than they look. On their legal status: "As administrative interpretations of the law, the Guides themselves are not enforceable." On their currency: the Commission opened a review — "With the present document, the Commission commences a new review" — noting that it "completed its most recent Green Guides review a decade ago," and asking whether to "initiate a proceeding to consider a rulemaking relating to environmental benefit claims under its FTC Act authority." That notice is dated the twentieth of December, 2022, with comments due the twenty-first of February, 2023.443 As of my writing, the Commission's own page on the subject still says the Guides "were first issued in 1992 and were revised in 1996, 1998, and 2012."444

One rule is moving, and it deserves credit. The Science Based Targets initiative's Corporate Net-Zero Standard, version 2.0, carries an effective date of the first of February, 2027, and contains a requirement that would have made much of the last decade's reporting impossible: "Companies shall determine target ambition based on the physical GHG inventory, including the location-based method for scope 2." It also excludes a list of mechanisms from counting toward target implementation, including "emissions reductions derived from hypothetical scenarios," credits "sold or otherwise transferred to a third party," and "emissions reductions or removals that are estimated ex ante and not yet realized."445 Ambition set on the physical grid. Not on the portfolio. That is the right rule, it is written down, and it takes effect in 2027 — by which time, on the schedules in section V, the plants will be under construction.

V.

The physical commitments, then. Nationally, the balance for the current year is not dramatic. The Energy Information Administration expects developers to add "86 gigawatts (GW) of new utility-scale electric generating capacity to the U.S. power grid in 2026," "a record if realized," of which solar is 51 percent, storage 28 percent, and wind 14 percent. Gas is a small slice: "In 2026, developers plan to add 6.3 GW of new natural gas-fired capacity," split "3.3 GW" of combined cycle and "2.8 GW" of combustion turbines.446 Against that, "4.6 GW of U.S. natural gas-fired capacity, almost 1% of the operating natural gas fleet at the end of 2025, is scheduled for retirement."447 Six point three in, four point six out. On the national ledger, gas is roughly treading water while solar and storage do the growing. But a national net figure is the wrong instrument. What matters is the duration of the new asset and the geography of the new load. Duration first. The federal model that projects the power system evaluates plant investment on a thirty-year horizon: "Technologies are compared on the basis of total capital and operating costs incurred over a 30-year period," and each option is judged "on the basis of a lifecycle cost over 30 years," with capital costs "recovered over the economic life of the asset."448 The same assumption tables give a combined-cycle construction lead time of three years and plant sizes of 418 and 1,083 megawatts.449 Three years to build. Thirty years to pay for. A unit ordered in 2026 is a revenue-recovering asset into the late 2050s — past the state net-zero dates, past the federal ones, past the budget in section

VI.

Now geography. Two states have written the data-center rationale into the record explicitly. A CITIZEN'S MANIFESTO FOR THE AGE OF AI 163 Georgia's Public Service Commission publishes a fact sheet that traces its own forecast. "In 2022, Georgia Power estimated it would need to increase power generation by only

400 megawatts over the next seven years." "By 2023, with the proliferation of data centers, that estimate grew to 6,600 MW." "Two years later, the estimate grew to 8,500 MW." And then the decision: "On December 19, 2025, the Commissioners voted to approve an agreement between the PSC's Public Interest Advocacy Staff and Georgia Power that certifies the construction of 9,985 MW of new energy generation – approximately 80 percent of which is expected to power data centers." For scale, the same document states that "Georgia Power currently deploys 22,000 MW statewide."450 The commission's own advisory on the intermediate step describes approving "at least 6,000 new megawatts of energy between 2029 and 2031," with a possible "additional 2,500 MW," and quotes its chairman: "As data center construction continues in Georgia, this IRP puts us in a safe and secure spot to meet that energy need."451 Roughly forty-five percent of what the utility currently deploys statewide, added in one authorization, four-fifths of it expected to serve one customer class. Louisiana wrote it into an order. On the twenty-ninth of August, 2025, the Public Service Commission found that construction of "the three new combined cycle combustion turbine ('CCCT') generators described in the Company's Application, testimony, and exhibits… serves the public convenience and necessity and is in the public interest, and is therefore prudent." The customer is named: "a data center being developed by Meta Platforms, Inc., through its subsidiary, Laidley LLC (the 'Customer'), in Richland Parish, Louisiana." The siting is stated: "Two of the CCCT generators will be constructed in Richland Parish. The third CCCT generator will be constructed in Killona, Louisiana."452 In the related record, the unit size and the contract term: "Each of the seven new CCCTs is expected to have a nameplate capacity of 754 MW," and the customer "will also contribute more than $1 billion toward ELL's embedded cost of service" "over the twenty-year Original Term of the ESA."453 Seven units at 754 megawatts is roughly 5,280 megawatts — that multiplication is mine — and the ground is already broken: "The Company broke ground in December 2025 on the Franklin Farms Power Station Units located in Richland Parish," with the third unit, "Waterford 5," expected to break ground in 2026. The record also describes the class in general terms: "data center customers like Laidley and Evest have load requirements exceeding 1GW, a high-capacity factor, and a requirement for 'speed to market.'"454

I have to state a limit here. I could not find, in the documents I fetched, the assumed book life or depreciable life of those combined-cycle units. The twenty-year contract term is in the record. The plant's life is not. That gap matters — a twenty-year agreement against a thirty- or forty-year asset is precisely the mismatch this chapter is about — and I am not going to fill it with an assumption. Ask for it. It exists in the depreciation schedules.

VII.

The reason I treat this as lock-in rather than as a phase is that the machines are already sold. GE Vernova reported in July 2026 that its "Gas Power equipment backlog and slot reservation agreements grew from 100 to 116 GW," that it "now anticipate[s] reaching at least 125 GW by year-end 2026," and that it remains "on track to deliver 20 GW of annual gas turbine output in the third quarter of 2026, with 24 GW in 2028," and is "implementing actions to produce 30 GW in 2030."455 One hundred sixteen gigawatts of book against twenty to twenty-four gigawatts a year of output is roughly five years of production already committed, on a capacity roadmap the company itself extends to 2030. Siemens Energy's own investor material states a new-unit backlog of about "78 GW" and, in three words, the market condition: "Capacity sold-out until FY28." The same deck names the demand: "24/7 reliable power critical to power data centers."456 Mitsubishi Heavy Industries reports a large-frame gas turbine backlog of eighty units, about thirty-five gigawatts, up from fifty-three units and twenty-three gigawatts a year earlier, with order intake "particularly strong in North America and Asia."457 Add the three books and the number is well over two hundred gigawatts of turbines contracted, from manufacturers

who describe themselves as sold out for years, to be installed in the late 2020s and early 2030s, financed on the thirty-year horizon the federal model assumes. That is the answer to the argument that this is a bridge. A bridge is a structure you cross and leave. These are assets whose cost recovery ends in the 2050s and whose physical lives may run longer. The decision about the 2050s is being made now, by purchase order, in a market where the queue for delivery is longer than the interval before the accounting standard changes. V

VIII.

The buildout's emissions are also not mostly where people look.

Return to Microsoft's data sheet. Scope 3 — everything upstream and downstream, which for a hyperscaler means chips, steel, cement, servers, and construction — was 18,243,000 tonnes of carbon dioxide equivalent in fiscal 2025, against a Scope 1 of 170,887 and a market-based Scope 2 of 2,707,428.437 Scope 3 is roughly six-sevenths of the total. That share calculation is mine; the components are the company's. Google reports the same asymmetry: operational emissions down two percent, "supply chain emissions grew by 25% year-over-year."434 A power purchase agreement, however well designed, however hourly, however local, addresses Scope 2. It does not touch the cement in the pad, the steel in the frame, the water and energy in the fabrication of the accelerators, or the diesel in the eighteen months of construction traffic. The part of the footprint most amenable to a contract is the smaller part. On the fabrication end, the largest contract manufacturer of advanced chips has published targets and a base year. Using 2025 as a baseline, the company "commits to achieving absolute reduction targets for scope 1, 2, and 3 emissions aligned with the SBTi Corporate Net-Zero Standard within the next decade," with milestones including "peaking carbon emissions in 2025, reducing them to 2020 levels by 2030 with a target of using 60% renewable energy for global operations (RE60), achieving RE100 by 2040, and reaching net-zero emissions by 2050." Its renewable share "surpassed 14% in 2024," and it has signed "a 20-year joint procurement agreement for 20,000 GWh of renewable energy."458 Peaking in 2025 and returning to 2020 levels by 2030 is a real commitment and I will say so. It is also an admission that the manufacturing layer's emissions are still climbing, and that fourteen percent renewable is the starting point.

IX.

There is a mismatch of clocks inside the buildout that I did not expect to find in a financial statement. Alphabet disclosed in its annual report that "In January 2023, we completed an assessment of the useful lives of our servers and network equipment and adjusted the estimated useful life of our servers from four years to six years," and that the effect "was a reduction in depreciation expense of $3.9 billion and an increase in net income of $3.0 billion" for the year.459 Amazon went the other way. "Effective January 1, 2024, we changed our estimate of the useful lives for our servers from five to six years, and effective January 1, 2025, we changed our estimate of the useful lives of a subset of our servers and networking equipment from six to five years." The reason: "The shorter useful lives are due to the increased pace of technology development, particularly in the area of artificial intelligence and machine learning." The effect: "an increase in depreciation and amortization expense of $1.4 billion and a reduction in net income of $1.0 billion."460

The same physical asset class has been assigned a four-year life, a five-year life, and a six-year life by different filers in different years, with billions of dollars of reported earnings turning on the choice—and one of the largest operators has now shortened the life, in writing, because of the pace of AI. Set that against the thirty-year cost-recovery horizon for the generating plants built to serve it.

Compute turns over about every half-decade. The generating plants built to serve it recover their capital over thirty years. We are financing thirty-year infrastructure to serve a five-year asset, on the assumption that the five-year asset will be replaced by something that needs at least as much power. And the retired hardware goes somewhere.

Compute turns over about every half-decade. The generating plants built to serve it recover their capital over thirty years. We are financing thirty-year infrastructure to serve a five-year asset, on the assumption that the five-year asset will be replaced by something that needs at least as much power. And the retired hardware goes somewhere. The most recent global accounting reports that "in 2022, a record 62 billion kg of e-waste was generated globally (equivalent to an average of 7.8 kg per capita per year); 22.3 per cent of this e-waste mass was documented as formally collected and recycled in an environmentally sound manner," and projects "82 billion kg of e-waste will be generated in 2030," with formal recycling rates declining "to 20 per cent in 2030" under business as usual.461 Twenty-two percent, falling. Against that baseline, Microsoft's self-reported "92% reuse and recycling of decommissioned servers and components for the second consecutive year"438 is genuinely good, and it is also the reason I want it in a permit rather than a report. A ninety-two percent rate that exists because a company chose it is a courtesy. The same number in a decommissioning condition is a requirement.

X.

Now the part that is personal, and that this book began with.

The federal interagency sea level rise assessment states the national picture: "Relative sea level along the contiguous U.S. (CONUS) coastline is expected to rise on average as much over the next 30 years (0.25–0.30 m over 2020–2050) as it has over the last 100 years (1920–2020)." For the Gulf it is worse: "For the observation-based extrapolations, the largest estimates of sea level rise in 2050 are found along the entire Gulf Coast. The Western Gulf has the highest extrapolated values in 2050, driven by high rates of coastal subsidence in the region." The Western Gulf figure for 2050 is 0.59 metres — about one foot eleven inches — with a likely range of 0.51 to 0.67 metres. The same report converts that into events: minor high-tide flooding is "projected to increase from a U.S average frequency of about 3 events/year in 2020 to >10 events/year by 2050," and moderate, "typically damaging" flooding "from a U.S. average frequency of 0.3 events/year in 2020 to about 4 events/year in 2050." And it contains one sentence about the federal flood maps that everyone building anything on a coast should read: "FEMA provides this information for national flood insurance purposes but does not consider future sea levels."462 The maps used to decide where it is safe to build are, by the acknowledgment of the agency that publishes the sea level projections, backward-looking. The regulatory floodplain is defined as "the area that will be inundated by the flood event having a 1-percent chance of being equaled or exceeded in any given year," "also referred to as the base flood or 100-year flood."463 A one-percent annual chance, computed from a record that stops before the trend. And under southern Louisiana, the land is going down. The state's own technical compilation reports geodetic subsidence rates "ranging from 0.7 to 7.1 mm/yr" across three basins, median shallow subsidence of "6.0 mm/yr for the Mississippi Delta and 5.8 mm/yr for the Chenier Plain," and satellite interferometry results for the New Orleans area showing 15 to 30 millimetres a year at Michoud and 10 to 20 in the Upper Ninth Ward and Bayou Bienvenue.464 Thirty millimetres a year is a foot a decade. I am not going to claim that a specific facility will flood, and I want to be exact about the claim I am making instead. One of the three combined-cycle units approved to serve the Richland Parish campus is sited at Killona, in the lower Mississippi delta, as Waterford 5.452 454 I do not have a hydrology study for that site and I am not asserting one. What I am asserting is narrower: the decision to build

was made using a flood map that its own publisher says excludes future sea level, in a state whose own documents record subsidence measured in centimetres per year, for an asset whose costs are recovered over three decades. That is not a prediction. It is a description of the instrument used to decide.

XI.

The grid is being asked to absorb this faster than its rules can be rewritten, and the federal regulator has said so three times in eighteen months. In November 2024 the Federal Energy Regulatory Commission rejected an amended interconnection agreement for a large load co-located at a nuclear plant. "In this order, we reject the Amended ISA," it wrote, finding that the regional transmission organization "has not demonstrated that the proposed non-conforming provisions in the Amended ISA are necessary deviations from the pro forma ISA due to specific reliability concerns, novel legal issues, or other unique factors." The chairman dissented: the agreement "represents a 'first of its kind' co-located load configuration that presents precisely the sort of specific reliability concerns, novel legal issues, and other unique factors that should have justified the filing of a non-conforming interconnection agreement." A concurring commissioner emphasized that the filing was "rejected without prejudice."465 The Commission then opened a general proceeding: "The Commission finds that PJM's tariff does not appear to sufficiently address the rates, terms and conditions of service that apply to co-location arrangements," directing the region and its transmission owners to "tell the Commission why the PJM tariff remains just and reasonable or, alternatively, what changes to the PJM tariff would remedy the Commission's identified concerns."466 At its December 2025 open meeting the Commission took up the result. I have to label this one carefully: what I was able to retrieve is the Commission's own summary of the draft order presented at the meeting, not the issued order itself. That summary states that "the draft order finds that PJM's Tariff is unjust and unreasonable because it does not contain provisions addressing with sufficient clarity or consistency the rates, terms, and conditions of service that apply to: Interconnection Customers serving Co-Located Load and Eligible Customers taking transmission service on behalf of Co-Located Load," and that "the Behind the Meter Generation (BTMG) rules in PJM's Tariff are no longer just and reasonable."467 I could not retrieve the issued order on the Commission's own site, and I am not going to quote its holdings from a mirror. What I can quote is the sequel: in June 2026 the Commission instituted a parallel proceeding for a second regional market, preliminarily finding that market's tariff lacks "sufficiently clear and consistent provisions" for integrating large and co-located loads, and defining the class it means — a new commercial or industrial customer at a single site with peak load of fifty megawatts or greater, interconnecting above sixty-nine kilovolts.468

Fifty megawatts is now a federal threshold for "this is a different kind of customer." The class was defined after the largest members of it were already under construction. The reliability picture is in the grid monitor's annual assessment, published in January 2026. "In the 2025 LTRA, NERC finds that 13 of 23 assessment areas face resource adequacy challenges over the next 10 years." Demand: "aggregated assessment area summer peak demand is forecast to rise by over 224 GW," which "is 69% higher than last year's 10-year growth projection of 132 GW," with winter peak growth "over 245 GW." Cause, stated without hedging: "New data centers for artificial intelligence and the digital economy account for most of the projected increase in North American electricity demand over the next 10 years." And the mismatch: "Projections for resource and transmission growth lag what is needed to support new data centers and other large loads that drive escalating demand forecasts." Meanwhile "the confirmed and announced-potential retirements over the next 10 years total over 105 GW in peak seasonal capacity."469 A ten-year demand forecast that rose by sixty-nine percent in a single

year is not a forecast. It is an admission that the planning system has lost the thread. And notice the climate consequence hidden in the reliability language: when demand growth outruns clean supply and retirements are pending, the balancing resource is the existing fossil fleet, run harder. That is the four percent rise in power-sector carbon dioxide from section II, arriving by arithmetic rather than by anyone's decision.

XII.

Which brings me home, to the case that made me start writing this book. Louisiana budgets for a coast. Its master plan is organized around two twenty-five-billion-dollar halves — "$25B RESTORATION" and "$25B RISK REDUCTION" — with candidate projects "evaluated based upon how much land they create and maintain over 50 years compared to a future without the master plan," and an outcome stated honestly: "Land area built or maintained that would have otherwise been lost by Year 50 ranges from 314 square miles in the lower scenario to 233 square miles in the higher scenario."470

The state's own climate plan puts the residual loss in a single sentence: "Even with a $50 billion investment in implementation of every project in the Coastal Master Plan, between 1,454 and 2,965 square miles of coastline are still likely to be submerged due to continued subsidence and sea level rise." The same document records the emissions commitment — the executive order "committed Louisiana to achieving net zero greenhouse gas (GHG) emissions by 2050" — the interim target — "reaching the interim 2030 goal of 40-50% reduction will be difficult" — and the structural fact that makes both difficult: "In the U.S., 17% of overall national emissions come from the industrial sector; in Louisiana, the industrial sector contributes an overwhelming 66% of overall state emissions."471 And the same state, through its economic development arm, announces that "Meta announced it is committing more than $50 billion to its Richland Parish Data Center project, expanding the campus to nearly 10 million square feet and 5 gigawatts of IT capacity," notes that "since 2024, Louisiana has secured over $150 billion in investment," and describes the energy arrangement: "The companies' recent agreement will fund seven new natural gas-fueled generating plants, three grid-scale batteries, potential nuclear development, including a project to increase the output of Entergy Louisiana's Waterford 3 nuclear facility and other purchased power." The same release records the offsetting commitments: "$215 million to its The Power to Care bill assistance program and residential energy efficiency programs over the course of the agreement and committing to fund up to 2.5 GW of clean, renewable energy."472 And the pipeline continues: the utility's own filing notes a ten-billion-dollar campus announced in West Feliciana Parish in December 2025 and twelve billion dollars across three campuses in Caddo and Bossier Parishes announced in February 2026.454 I want to be scrupulous about two things. First, the offsetting commitments are real and I am not sneering at them. Two hundred fifteen million dollars of bill assistance is money in the hands of people who need it. Up to two and a half gigawatts of renewables is a large procurement. A nuclear uprate is genuinely low-carbon generation. Second, my strongest claim here is an absence, not a contradiction: no document I fetched from the State of Louisiana reconciles the three commitments. There is a restoration plan that assumes a rising, subsiding Gulf. There is a climate plan with a 2030 interim target and a 2050 net-zero target. There is an industrial recruitment program adding seven gas-fired units for one customer class. Each is defensible on its own paper. I could not find the paper on which they are added together.

I also cannot tell you the current legal force of that climate plan. It was submitted in February 2022 under a previous administration, and I could not verify whether it has been reaffirmed, superseded, or quietly abandoned. If it has been abandoned, then the reconciliation problem does not disappear — it simply resolves

in the direction of the buildout, and the honest thing would be to say so out loud, in a document, so that people can vote on it. That is the whole of my complaint about the cumulative question, and it is not a complaint about gas or about data centers. It is that three arms of one government are making three commitments with the same dollars and the same atmosphere and the same coast, and no instrument anywhere requires them to be added up.

XII.

The concerns, and where each is actually decided. If you get honest answers to the middle column, you do not need me. If you cannot get them, that is the finding, and it is reportable on its own.

XIII.

The three tests, applied to this concentration. Witness. The witness problem in this chapter is not secrecy. It is asymmetry of arrival. Every document I have quoted is public: the regulation, the standard, the order, the fact sheet, the assessment, the data sheet. The honest measure arrives after the decision. The hourly rule takes effect in 2030 for hydrogen and in 2027 at the earliest for corporate accounting, with a legacy clause for contracts already signed.425 427 The net-zero standard's requirement to set ambition on the physical grid takes effect in February 2027.445 The federal definition of a fifty-megawatt large load arrived in 2026, after the largest such loads were under construction.468 A witness admitted after the verdict is a spectator. Burden. The burden here has an unusual shape, and I want to name it precisely. It is not principally a bill; the previous chapter dealt with bills. It is a commitment made on someone else's behalf across time. A thirty-year cost-recovery period is a claim on the electricity system of the 2050s, entered into in 2026 by a commission, a utility, and a customer, none of whom will be in the room in the 2050s.448 A twenty-year customer agreement against an asset of unstated longer life leaves a remainder, and remainders are paid by whoever is still connected.453 The four percent rise in power-sector emissions is borne, in the first instance, by a hot summer's asthma admissions and, in the last instance, by a coast.431

Almost nothing in this chapter is designed to be returned, and one thing is. The nuclear uprate adds low-carbon output to a grid that needs it.472 The recycling rate of ninety-two percent is a genuine reclamation of material.438 The twenty-year renewable procurement in the chip supply chain is real capacity built.458 Every one of those is voluntary, and therefore revocable. Meanwhile the things that cannot be returned — a carbon budget spent, a delta that has subsided, a thirty-year asset financed — are the ones no instrument requires anyone to guarantee. That is the inversion this book keeps finding: the reversible commitments are enforceable, and the irreversible ones are not.

XIV.

Here is what I would insist on. Six requirements. Each names who decides it and what happens if they do not.

1. Dual disclosure of both numbers, in the same document, by any operator over fifty megawatts. The location-based Scope 2 figure and the market-based figure, published side by side as the accounting standard already requires, plus the hourly matched percentage for each grid region in which the operator draws power. Both numbers, on the same page, before the certificate is granted.440 441 Who: the state utility commission, as a condition of the service agreement; the securities regulator for public filers; the state attorney general under

consumer-protection law. If not: the operator may not make any public claim about clean or renewable energy in that jurisdiction, and any incentive conditioned on such a claim is void.

2. No new fossil generation for a single private load without a bonded retirement date. If a combined-cycle unit is certified to serve an identified customer, the order approving it must state the retirement year and require security sufficient to cover early retirement, sized in the same proceeding. The customer agreement's term and the plant's cost-recovery period must be stated as two numbers with the gap between them written out.448 453 Who: the utility commission that issues the certificate; the legislature where the commission lacks authority to require security. If not: the certificate does not issue, and cost recovery for the plant is presumptively denied in the next rate case.

3. Hourly, locational matching as the standard for public benefit — not for accounting alone. Any operator receiving a tax abatement, a discounted rate, an expedited permit, or a public incentive must demonstrate matching on the hourly, same-grid basis the government already requires of hydrogen producers, on a stated phase-in, with the incentive tapering as the percentage falls short.425 426 Who: the local governing body granting the abatement; the state revenue department; the legislature that authorizes the exemption. If not: the abatement converts to a payment in lieu, at the full assessed value, for the years the standard was not met.

4. Lifecycle emissions in the record, including construction and hardware. The application must include the embodied emissions of the build — concrete, steel, servers, accelerators — and the projected refresh interval with the reuse and recycling rate for decommissioned equipment, along with its destination.437 460 461 Who: the county or parish issuing the land-use approval; the state environmental agency. If not: the application is incomplete, and the clock on any deemed-approval provision does not start.

5. Siting decided against projected exposure, not the current insurance map. Any facility or generating unit above a stated size, in a coastal, subsiding, or floodplain parish, must be evaluated against the published sea level and subsidence projections for its own region, with the storm contingency plan for fuel, coolant, and workers filed publicly. The federal flood map is a floor, and its publisher says it excludes the future.462 463 464 Who: the parish or county permitting authority; the state coastal agency; the utility commission for generating units. If not: no permit, and no cost recovery for damage that the projection anticipated.

6. One cumulative ledger per region, published before the next approval. A single document, maintained by the state, adding up every approved and pending large load in the region: megawatts, water, acres, new fossil capacity, and projected emissions, set against the state's own emissions target and its own coastal plan. A running total, updated before each new approval, that anyone can read.470 471 472 Who: the governor's office or the state environmental agency, by executive order; the legislature, by statute, if the executive will not. If not: no further certificates for loads above the threshold until the ledger is published and current. If you can carry only two, carry the first and the sixth. Everything else in this chapter is an argument about how to allocate a burden. Those two are about whether the burden is ever written down in one place, in numbers, before it is incurred — and no one who has read the documents I have quoted can tell you that it currently is.

XV.

The accounting problem is not that annual renewable-energy certificates are necessarily false or illegitimate. A company may purchase certificates equal to its annual electricity consumption, retire them properly, and

accurately state that it has matched its electricity use on an annual volumetric basis. That is the established accounting instrument, and it has helped create demand for renewable generation.

But it is not a physical description of what powered a facility in a particular place at a particular hour. A company can accurately report a one-hundred-percent annual renewable match while drawing electricity, during its highest-load hours, from a grid whose marginal generation is substantially fossil-fueled. The annual number is true within the rules that produce it. The hourly, location-based number may tell a materially different story. Until recently, no general rule required companies to publish both.

The distinction becomes visible when firms choose the harder disclosure. One company published an hourly carbon-free-energy percentage in the sixties alongside a one-hundred-percent annual renewable-energy claim. Another saw its reported emissions rise sharply after ending its use of renewable-energy certificates it characterized as non-additional. Neither disclosure establishes that the company is failing; both establish that annual matching and physical, hourly electricity use are different measures and should not be presented as interchangeable.

That is why the relevant question whenever a data center is said to run on clean energy is not simply whether certificates were purchased and retired. Ask which number is being claimed: the annual market-based match, the location-based grid figure, or the hourly, deliverable clean-energy percentage. Ask for all three, and ask which hours and which grid the claim covers. The people making the claim are not necessarily lying. Most are using the accounting rule available to them. The problem is that the rule was designed for a slower buildout and is being revised too late to govern much of the infrastructure now being financed and constructed.

Are we spending the remaining carbon budget on the tool we claim will solve the carbon problem? No one can answer that responsibly from the present record. The range of projected load is too wide, the accounting rule is still being rewritten, and the relevant figures are not required to appear in one place before the concrete is poured. But the uncertainty is not an excuse for proceeding blindly. It is the reason to require the ledger before the next approval. That is a choice, it is recent, and it can be changed by a state legislature in a single session. — D.M.D.

Chapter Nine

Chapter Nine: Who Owns the Machine

Concentration, Capital, and the Four-Company Problem

If a handful of firms own the models, the chips, the clouds, and the distribution, in what sense is any of this a market?

I.

I want to begin with a single document, because a single document settles the question faster than any argument I could make. On the twenty-fifth of January, 2026, NVIDIA closed a fiscal year and filed the annual report that securities law requires it to file. The company reported $215,938 million in revenue, of which $193,737 million came from its data center segment.473 Then, in the section where a filer must disclose its dependence on particular buyers, it wrote this: "For fiscal year 2026, sales to one direct customer represented 22% of total revenue and sales to another direct customer represented 14% of total revenue."474 Two buyers, thirty-six percent of the revenue of the most valuable company in the world. Elsewhere in the same filing, the company disclosed what it had done with its cash. "We invested $17.5 billion in private companies and infrastructure funds, primarily to support early-stage startups. These investments include AI model makers that purchase our products directly or through CSPs."475 Its holdings of non-marketable equity securities rose from $3,387 million to $22,251 million in twelve months.476 Elsewhere again, it disclosed what it had promised to buy. "Multi-year cloud service agreement commitments as of January 25, 2026, were $27 billion" — cloud capacity purchased from the very cloud providers that are its largest customers, with the caution that "[s]ome cloud service capacity may be reduced, terminated or sold to others by the CSPs, in which case our commitments will be reduced."477 It reported $95.2 billion in purchase obligations, $11.4 billion in investment commitments, and "$3.5 billion in land, power, and shell guarantees to early-stage companies."478 And then, in a sentence I have read perhaps thirty times and still find remarkable, it disclosed a revenue channel it could not or would not name: "We estimate that one AI research and deployment company contributed to a meaningful amount of our revenue purchasing cloud services from our customers in fiscal year 2026."479

Read those disclosures together. The chip supplier is an equity investor in the model developers. The model developers are its customers. The cloud providers are also its customers, and it is simultaneously their customer, having committed twenty-seven billion dollars to buy back capacity from them. An unnamed model developer generates meaningful revenue for the chip supplier by renting the chip supplier's chips from the chip supplier's customers. Money leaves the supplier as investment, returns as revenue, and is counted as growth at both ends. I am not alleging fraud. Every one of those sentences was written by the company's own lawyers, audited, and filed with the Securities and Exchange Commission, which is precisely my point. This is not a hidden structure. It is the disclosed structure. A market is a place where independent buyers and independent sellers discover a price. This is a place where the same capital appears on four sides of the same transaction and calls the result demand.

II.

Let me also show you what the same filing does with something the company had already announced. On the twenty-second of September, 2025, NVIDIA issued a press release describing a letter of intent with OpenAI, framed around an investment of "up to $100 billion." Five months later, in the audited annual report, the company wrote: "We are finalizing an investment and partnership agreement with OpenAI. There is no assurance that we will enter into an investment and partnership agreement with OpenAI or that a transaction will be completed."480 That sentence appears twice in the filing. Citizens should learn to read exactly this gap. The press release is marketing and carries no liability. The 10-K carries liability, and so it tells the truth: the hundred billion dollars was a headline, not yet a contract. When you are told that a hundred billion dollars is flowing into your state, your grid, or your industry, the number to ask for is not the one in the announcement. It is the one in the filing. Two quarters later the filing had something firmer to say. In its quarterly report for the period ended the twenty-sixth of July, 2026, NVIDIA disclosed that "[i]n August 2026, we entered into guarantees, capped at a total of $105 billion, to provide credit support on a land, power, and shell buildout with affiliates of SB Energy Corp. (SB Energy) on behalf of a customer, an affiliate of OpenAI Group PBC (OpenAI), related to leases for approximately 4.25 gigawatts of IT load in the aggregate at SB Energy's PORTS Technology Campus in Pike County, Ohio."481 Nine phases. Twenty-year lease terms. First phase expected in fiscal 2029. An option on roughly 3.8 additional gigawatts. And the consideration: "In exchange for the guarantees, the site will exclusively host NVIDIA AI infrastructure, subject to limited exceptions."481

Hold that word. Exclusively. A chip vendor has underwritten a hundred and five billion dollars of a customer's real-estate and power obligations, and in return four gigawatts of American industrial capacity is contractually closed to its competitors. By the end of that quarter the company held $47,898 million in privately held equity securities and disclosed equity investments of $99 billion with a further $25 billion committed.482 That is not a supplier. That is a landlord, a lender, a shareholder, and a zoning authority operating by contract.

III.

Underneath all of it sit two companies most citizens have never heard of. ASML, a Dutch firm, states plainly in its own annual report that "ASML is currently the world's only manufacturer of EUV lithography systems," and it recognized revenue on forty-eight of those systems in 2025.483 Every leading-edge logic chip on earth is patterned by a machine made by one company in one country. ASML in turn depends on Carl Zeiss SMT for critical optics, and states that were that supply interrupted, "we would effectively cease to be able to conduct our business."484 Downstream, TSMC reports that advanced technologies — seven nanometres and below — accounted for seventy-four percent of its wafer revenue, and that its top ten customers accounted for seventy-eight percent of net revenue in 2025.485 So the physical base of the whole edifice is: one optics maker, one lithography maker, one dominant advanced foundry, one dominant accelerator designer, and roughly four buyers who matter. Each link is a company with real engineers doing genuinely hard things. I do not begrudge any of them their competence. I object to the arithmetic. Five single points of failure in series is not an industry. It is a supply line, and supply lines are what you build when you are preparing to be told what to do.

IV.

Above the silicon, the coupling of cloud and model is now nearly complete, and here the interesting evidence is often what a filing declines to say. DAVID M. DOYLE 178 Microsoft's annual report describes an

"approximate 25 % interest" in OpenAI Group PBC, accounts for it under the hypothetical liquidation at book value method, records funding of $13.0 billion of a $24.1 billion commitment, and discloses "rights to OpenAI" intellectual property.486 Microsoft's own blog post describing the same restructuring called it "roughly 27 percent."487 Two percentage points of a company at that valuation is tens of billions of dollars. I take the audited number over the blog post, and I note for citizens that when a firm gives you two figures for the same thing, the one it publishes voluntarily is the one designed to be read. Amazon reports its Anthropic position through convertible notes and equity, alongside a large committed cloud spend. Alphabet, likewise an investor in Anthropic, presents a starker case. I searched the full text of Alphabet's most recent annual report and its most recent quarterly report for the word "Anthropic." It appears zero times in either document.488 One of the two or three most consequential relationships in the industry — a multibillion-dollar stake in a leading model developer that is also a major purchaser of Google Cloud — is, in the judgment of Alphabet's disclosure counsel, not material enough to name. The Federal Trade Commission looked at exactly these three arrangements under its Section 6(b) study authority and published staff findings in January 2025. The report examined "three partnerships involving generative artificial intelligence ('AI'): Microsoft-OpenAI, Amazon-Anthropic, and Google-Anthropic."489 Its findings, in staff's own words: "The partnerships provide CSP partners significant equity and certain revenue-sharing rights in their AI developer partners. The partnerships also leave open the possibility that the AI developers could be fully acquired by the CSPs in the future."490 "In some cases, the agreements include exclusive or preferential treatment of the CSP partner and the ability to potentially influence the AI model developer partner through board seats and consultation rights."491 And then the sentence that names the thing: "The partnerships include cloud commitments that require AI developers to spend a large portion of their CSP partner's investment on cloud services from their partner. This feature of the partnerships—circular spending—raises questions about the extent to which the investments represent genuine capital infusions versus a mechanism to lock in cloud business."492 Staff also found that the agreements "may increase switching costs via contractual commitments (including exclusivity requirements) and technical barriers," that Microsoft built a supercomputer "in collaboration with and exclusively for OpenAI," and that OpenAI had used Microsoft as its "exclusive cloud provider" from 2019 onward.493 The report quotes Microsoft's chief executive on what the arrangement secures without ownership: "If OpenAI disappeared tomorrow … we have all of the rights to continue the innovation."494

An agency of the United States government identified circular spending, exclusivity, board influence, switching costs, and an acquisition option, described them accurately, and then — as far as any public record I could locate shows — that was the end of it. No enforcement action followed. The report's own currency note is worth keeping: its findings "represent information available to staff as of September 2024 as well as publicly available information through January 2025."495 Everything in this chapter's first two sections happened after that date.

V.

The circularity is not confined to the frontier labs. CoreWeave, the specialist cloud provider, discloses that it derived sixty-seven percent of revenue from its top customer, that "NVIDIA Corporation invested" in the company, and that "all of the GPUs used in our infrastructure" come from NVIDIA.496 One supplier, one customer, one vendor who is also a shareholder. Oracle presents the same structure in the negative. Its annual report states that "[n]o single customer accounted for 10% or more of our total revenues," which is technically a statement about concentration and functionally a denial of it.497 The same filing reports remaining performance

obligations of $638 billion, up from roughly $138 billion a year earlier, and free cash flow of negative $23.7 billion.498 Its quarterly release explains part of the mechanism: "the customer prepaid Oracle for the purchase of the GPUs, or the customer bought and supplied the GPUs to Oracle. The prepaid and customer supplied hardware portions of our large AI contracts now total $75 billion."499 I searched the full text of Oracle's annual report for "OpenAI." Zero occurrences.497 Meta's Louisiana joint venture works the same way from the other end. Meta holds a twenty percent membership interest in a venture carrying roughly $27 billion of development costs, discloses "maximum exposure to loss related to the Venture was $45.95 billion," and does not consolidate the venture onto its balance sheet.500 A data center the size of a small city, financed off the books of the company for whom it is being built. I dwell on these because of what they do to accountability. When a facility is owned by a venture, financed by a pension fund, guaranteed by a chip vendor, powered under a confidential tariff, and operated for a tenant who is an equity affiliate of the guarantor, there is no single door for a citizen to knock on. Diffusion of ownership is not decentralization. It is the engineering of an absent defendant.

VI.

Then there is talent, which the law has never learned to see.

In 2025 Meta paid, by its own disclosure, $13.80 billion for a non-controlling interest in Scale AI, and stated in the same filing that it does "not have significant influence" over the company.501 Scale's own announcement described the transaction only as covering "a minority of Scale's outstanding equity" and named no percentage.502 The forty-nine percent figure that circulated in every news account of the deal appears, so far as I can find, in no company document from either side. The government has already studied the pattern. In September 2021 the Federal Trade Commission's chair described a staff review of "over 800 unreported acquisitions by the most significant players in digital markets—Alphabet, Amazon, Apple, Facebook, and Microsoft," noting that these firms had "devoted tremendous resources to acquiring start-ups, patent portfolios, and entire teams of technologists—and how they were able to do so largely outside of our purview."503 She drew the obvious conclusion: "this study underscores the need for us to closely examine reporting requirements under the Hart-Scott-Rodino Act and to identify areas where the FTC may have created loopholes that are unjustifiably enabling deals to fly under the radar."504 The study found that "over 76% of the acquisitions captured including non-compete clauses for founders and key employees," and that "[e]xploring how firms in digital markets may be using acquisitions to lock-up key assets along with talent will be a worthy area of study."505 Five years later I could not locate any public statement from either antitrust agency taking a position on whether the current generation of talent-and-license transactions is reportable at all. If it exists, it is not where a citizen would look. The loophole was identified, published, and left open, and the structure has since been used to move the leadership of entire model teams between competitors without a merger filing being made.

VII.

Concentration of ownership would matter less if the knowledge stayed public. It has not. The technical report accompanying GPT-4 states the policy without embarrassment: "Given both the competitive landscape and the safety implications of large-scale models like GPT-4, this report contains no further details about the architecture (including model size), hardware, training compute, dataset construction, training method, or similar."506

Two reasons are offered — competition and safety — and they are joined so that the second launders the first. Competitive advantage is a business interest; the public has no obligation to protect it. Safety is a public interest; the public has every obligation to weigh it. By binding them in one sentence, the disclosure that serves the shareholder is dressed in the language that serves the citizen. If the argument for secrecy is safety, then the details should go to a regulator under seal rather than to nobody. If the argument is competition, say so alone and let the public decide what it is owed. The consequence is that the most important industrial process of this decade is now less documented, in public, than the process for making cement. An independent scientist cannot replicate a frontier training run. A regulator cannot audit one. A citizen at a rate hearing cannot even establish how much computation the facility she is being asked to subsidize will actually perform.

VIII.

Into that documentary vacuum has stepped private government. The Frontier Model Forum states that "[t]he current members of the FMF are Amazon, Anthropic, Google, Meta, Microsoft, and OpenAI."507 Six firms — the same six that appear as investors, customers, cloud hosts, and counterparties throughout this chapter — convened to develop shared practice on frontier model safety. I do not doubt the sincerity of the people involved. I object to the constitutional form. A body composed exclusively of the regulated, producing the standards by which they will be judged, with no seat for anyone who bears the risk, is a trade association performing sovereignty. The self-binding documents say as much. Anthropic's responsible scaling policy concedes that "we cannot commit to following them unilaterally" and provides that the chief executive and the responsible scaling officer "will make the ultimate determination."508 OpenAI's preparedness framework establishes an internal Safety Advisory Group that sets and clears its own capability thresholds, and includes a provision contemplating adjustment of its commitments if a competitor ships a comparable system without equivalent safeguards.509 Those are not commitments. They are intentions with a named officer authorized to revoke them and a clause that makes the least careful competitor the effective legislator for all of them. In every other safety-critical industry we long ago concluded that this arrangement does not hold: pressure vessels, aircraft, pharmaceuticals, bridges. The rule we settled on is that the party bearing the risk sets the threshold and an outside body verifies compliance. There is no version of the AI story in which that principle is inapplicable. There is only a version in which it has not yet been applied.

IX.

Now to the state, and here I have to report a pattern rather than a failure, because the individual actors mostly did their jobs. In August 2024, after a full trial, Judge Amit Mehta of the District Court for the District of Columbia held that "Google is a monopolist, and it has acted as one to maintain its monopoly. It has violated Section 2 of the Sherman Act."510 The court found that Google's distribution agreements were exclusive and anticompetitive, that Google had charged supracompetitive prices for search text ads, and that it had earned monopoly profits.510 At remedies, the court described what those agreements had done: they "prevented rivals from accumulating the queries and associated data, or scale, to effectively compete," and "[t]aken together, these agreements effectively 'froze' the search ecosystem, resulting in markets in which Google has 'no true competitor.'"511 The access points covered "account for roughly 50% of all search queries issued in the United States," default placements drove "70% of all U.S.-based queries" at the time of trial, and in 2021 Google "paid more than $26 billion in 'traffic acquisition costs' to distribution partners."512 The court also understood what the case had become. "But artificial intelligence technologies, particularly generative AI ('GenAI'), may yet

prove to be game changers," it wrote, noting that "tens of millions of people use GenAI chatbots … to gather information that they previously sought through internet search."513 "The emergence of GenAI changed the course of this case. No witness at the liability trial testified that GenAI products posed a near-term threat to GSEs. The very first witness at the remedies hearing, by contrast, placed GenAI front and center as a nascent competitive threat."514 The proceedings were "as much about promoting competition among GSEs as ensuring that Google's dominance in search does not carry over into the GenAI space."515 Then came the remedy. Exclusivity was barred: "Google will be barred from entering or maintaining any exclusive contract relating to the distribution of Google Search, Chrome, Google Assistant, and the Gemini app."516 Bundled and multi-year conditioning was barred, and the final judgment requires that any agreement conditioning consideration on placement of a Google search or generative-AI product "terminates no more than one year after the date it is entered."517 The decree defines a "GenAI Product" to include any product whose principal functions include "answering information-seeking prompts across a wide variety of topics using a broad range of publicly available information," sweeps such products into the definition of a search access point, and forbids an exclusive Apple arrangement for any Google generative-AI product in Safari or on Apple devices.518 The extension of the decree to generative AI before the market had consolidated is the single most farsighted act by any organ of the United States government described in this chapter.

But the court declined to ban the payments. "Google will not be barred from making payments or offering other consideration to distribution partners for preloading or placement of Google Search, Chrome, or its GenAI products," it held, reasoning that cutting off the payments would impose "substantial—in some cases, crippling—downstream harms to distribution partners, related markets, and consumers."519 The court had stated the case against its own conclusion two pages earlier: a payment ban "would pry open the market to competition," because the revenue-share payments "'provide an incredibly strong incentive for the ecosystem to not do anything'" and their "'net effect [is to] basically freeze the ecosystem in place.'"520 And it acknowledged the cost: "The court well recognizes what eschewing a payment ban may mean for competition. Due to Google's massive financial advantage and its superior monetization, distributors will be incentivized to stick with Google because it can pay more, thus leaving in place the very forces that 'effectively [have made] the ecosystem exceptionally resist[ant] to change.'"521 It reserved the power to revisit: "The court is thus prepared to revisit a payment ban (or a lesser remedy) if competition is not substantially restored through the remedies the court does impose."522 I want to be careful here, because it would be cheap to call that timidity. It is not. It is a judge with a defined remedial mandate declining to inflict collateral damage on third parties who did not violate the law — Mozilla, handset makers, carriers — in order to reach a defendant who did. The reasoning is honest and the risk is stated on the record. My objection is not to the judge. It is to a system in which the remedy for a proven monopoly depends on how many other firms have been made financially dependent on it. Dependence, in that system, is not merely the harm. It is the defense.

X.

The same year produced two more decisions, and together they complete the picture.

In April 2025 Judge Leonie Brinkema of the Eastern District of Virginia held that "Plaintiffs have proven that Google has willfully engaged in a series of anticompetitive acts to acquire and maintain monopoly power in the publisher ad server and ad exchange markets for open-web display advertising," and that "Google is liable under Sections 1 and 2 of the Sherman Act."523 Sixteen months later, on the second of September, 2026, the court

ordered "that Plaintiffs' proposals for structural remedies in the form of the divestiture of AdX, the open-sourcing of DFP's final auction logic, and the contingent divestiture of DFP Remainder, be and are REJECTED; and it is further ORDERED that most of the parties' proposed behavioral remedies, as modified by this Court, be and are ACCEPTED."524 The reasoning was sealed for fourteen days for redaction, and as of this writing I could not obtain a public version of it.524 So I can tell you what was decided and not why, which is itself a fact about how much of this governance is legible to the governed. In December 2025 Chief Judge James Boasberg entered judgment for Meta in the FTC's monopolization case, concluding that "the agency has not carried its burden: Meta holds no monopoly in the relevant market."525 "While it once might have made sense to partition apps into separate markets of social networking and social media, that wall has since broken down."526 The court found that "personal social networking is not a separate product market," that Meta competes in a market including TikTok and YouTube, and that on that definition "Meta's share of that market comes out to around 30% of time spent."527 "Even if YouTube is out, including TikTok alone defeats the FTC's case."528 What broke down the wall? TikTok runs "an AI algorithm that shows users content based not on their friends, but on their interests," and what has replaced content from friends is, "[f]or the most part, short videos posted by strangers and recommended by AI."529 A recommendation model dissolved the market definition on which a decade of antitrust theory rested, and the dissolution of the category became the defense. The same technology that concentrates power at the infrastructure layer decomposes the legal categories that would measure power at the application layer. If our doctrine describes markets as bundles of features, and machine learning can recombine features at will, the doctrine will always be one product cycle behind the firms it governs.

XI.

That is the judiciary. The executive branch did something stranger: it built the one instrument that could have measured concentration, and then deleted it. A CITIZEN'S MANIFESTO FOR THE AGE OF AI 185 Executive Order 14110, signed in October 2023, invoked the Defense Production Act to require developers of dual-use foundation models to report to the federal government on their training activities, on "the ownership and possession of the model weights," and on red-team results; and to require anyone who "acquire[s], develop[s], or possess[es] a potential large-scale computing cluster to report any such acquisition, development, or possession, including the existence and location of these clusters and the amount of total computing power available in each cluster."530 It then drew the line: reporting applied to "any model that was trained using a quantity of computing power greater than 10^26 integer or floating-point operations," and to any cluster with "a theoretical maximum computing capacity of 10^20 integer or floating-point operations per second for training AI."531 The same order stated that federal policy required "addressing risks from dominant firms' use of key assets such as semiconductors, computing power, cloud storage, and data to disadvantage competitors."532 Whatever one thinks of that threshold, it was the only mechanism by which the United States government would have known, as a matter of routine record, where the compute in this country is, who owns it, and how much of it there is. Every question in this chapter would be easier to answer if that register existed. On the twentieth of January, 2025, Executive Order 14148 revoked it, in a list, at item "(ggg)."533 Three days later Executive Order 14179 directed that all actions taken under the revoked order be reviewed and, as appropriate, "suspend[ed], revise[d], or rescind[ed]," and set the replacement policy: "to sustain and enhance America's global AI dominance."534 America's AI Action Plan, published in July 2025, described the revoked order as one "that foreshadowed an onerous regulatory regime" and declared that "America's private sector must be

unencumbered by bureaucratic red tape."535 I searched the full text of that plan for the words "monopoly," "concentration," and "antitrust." The only near-monopoly it names is a favorable one, held by "America and its close allies" over parts of the semiconductor pipeline, and the competition it discusses throughout is competition with China, not competition within the United States.536 Domestic market structure is not addressed. It is not rejected as a concern. It simply does not appear.

Compare Europe, not because Europe is wiser but because the contrast isolates the variable. Under the Digital Markets Act, an undertaking is designated a gatekeeper on stated quantitative presumptions — turnover of €7.5 billion, or market capitalisation of €75 billion, plus forty-five million monthly active users and ten thousand business users in the Union — and gatekeepers are then barred from self-preferencing, required to provide effective interoperability free of charge, required to grant fair, reasonable and non-discriminatory access, and barred from combining personal data across services without specific consent.537 Twenty-three core platform services provided by seven gatekeepers are currently designated.538 None is an AI model or a cloud service: the designated categories remain search, social, intermediation, ads, browsers, operating systems, and messaging.539 Europe built a structural instrument and has not yet pointed it at this industry. On compute, though, Europe kept the line the United States dropped. Article 51 of the AI Act provides that "[a] general-purpose AI model shall be presumed to have high impact capabilities pursuant to paragraph 1, point (a), when the cumulative amount of computation used for its training measured in floating point operations is greater than 10^25," with the Commission empowered to move the threshold by delegated act as hardware improves.540 That figure is one order of magnitude below the American trigger, it is binding law rather than executive discretion, and it is in force. The jurisdiction hosting nearly all of the world's frontier training compute has no compute register, and the jurisdiction hosting comparatively little has one.

XII.

Meanwhile the government became a customer, on terms it would not have accepted from a defense contractor. Those are the General Services Administration's own published figures: "Every participating U.S. federal agency will have access for a nominal fee of $1 to ChatGPT Enterprise for one year"; agencies "paying $0.47 per agency for Google's AI tools"; Grok "for only $0.42 per organization."541 The agency reports "approximately $1.68 billion in cost savings," of which "approximately $1.4 billion stems from AI agreements that have expanded access to AI tools for roughly 3.5 million federal employees."542 A next-phase agreement, described as "a 27-month offer expected to take effect Oct. 1, 2026," moves to discounted consumption-based pricing.542 The agency's own logic is candid: "agencies wanted the ability to try new technologies before committing to a lengthy acquisition. As agencies increasingly integrate AI into their regular operations, providing consumption-based access is the next logical step."543

Set that beside the policy the Office of Management and Budget issued on the third of April, 2025. "As agencies seek to accelerate the adoption of AI-enabled services, they must pay careful attention to vendor sourcing, data portability, and long-term interoperability to avoid significant and costly dependencies on a single vendor."544 Agencies should seek proposals "that reduce the risk of vendor lock-in, such as requirements regarding knowledge transfer, clear data and model portability practices, clear licensing terms, and pricing transparency."545 Practices to be discouraged include "offering products without bulk pricing arrangements, tying arrangements, steering arrangements, minimum spend requirements, or other agreements that encourage consolidation of spending with one vendor or one group of vendors through fixed contract lengths, exclusive

discounts, or other incentives."546 Vendors should offer "systems or services at uniform and publicly available prices and not engag[e] in self-preferencing."547 A companion memorandum instructs agencies to prefer interoperable products and to retain rights in their own data and any improvements to it.548 I want to state the tension precisely, because I could not resolve it. A one-dollar, twelve-month price with a fixed expiry, extended to millions of employees who will build workflows on it, is the textbook shape of an introductory offer whose economics arrive at renewal. Whether the underlying agreements contain the portability, knowledge-transfer, and exit terms the memorandum requires, I cannot say: no agreement text, ordering guide, or contract clause is published on the agency's site, and I looked.549 The government has written the correct rule about dependency and simultaneously executed the transaction that most efficiently creates it, and it has not published the documents that would tell a citizen which instinct won. The honest version of the government's position is this. Agencies were being left behind, the tools were expensive, and someone found a way to get them into the hands of three and a half million public servants for the price of a coffee. That is a real accomplishment and I will not pretend otherwise. But an introductory price is a form of leverage, and leverage extended to a state is a constitutional matter, not a procurement one.

XIII.

The alternative — public compute — exists, and its scale tells you how seriously it is meant.

The National Artificial Intelligence Research Resource is described by the National Science Foundation as "a scalable national infrastructure that provides the research and education communities with access to the critical computing, software, data, models, educational resources and expertise needed to advance AI innovation," now "supporting more than 600 research projects and 6,000 students across all 50 states, Washington, D.C. and Puerto Rico."550 Read the donated allocations as published. Anthropic: "API access to Anthropic's Claude model for 10 researchers working on climate change-related projects."551 Amazon Web Services: "Support for at least 20 research projects through access to AWS credits."552 Groq: "access for up to 10 research teams."553 Hugging Face: "100 compute grants."554 Microsoft: "$20 million in compute credits."555 Hewlett Packard Enterprise, still hypothetical: "Proposed contribution: Time on GPU-powered supercomputing platforms and discounts."556 Meta's contribution is bounded by its own licence terms.557 Ten researchers. Twenty projects. Twenty million dollars in credits, against the ninety-five billion in purchase obligations disclosed by a single vendor in a single filing.478 This is not a public option. It is a philanthropic gesture from the firms whose market position the public option would exist to counterbalance, and it can be withdrawn at renewal. The government's own strategy document confirms the dependency by directing agencies to "[p]artner with leading technology companies to increase the research community's access to world-class private sector computing."558 The statutory fix has been introduced three times and enacted zero times. The CREATE AI Act of 2024 was placed on the Senate calendar in December 2024 and died with the Congress.559 A House version was ordered reported, twenty-nine to nothing, in June 2026.560 A Senate version was read twice and referred to committee in April 2026.561 Unanimous in committee, absent from the statute book. The nation's public research computing capacity therefore rests on agency discretion plus donated credits, which is to say it rests on the goodwill of the parties it was designed to check.

XIV.

One instrument did treat compute as a matter of sovereign geography, and it lasted four months.

The Bureau of Industry and Security's interim final rule on the diffusion of artificial intelligence, effective in January 2025, revised export controls on advanced computing chips and, for the first time, imposed controls "on the model weights of certain advanced closed-weight dual-use AI models controlled under newly created ECCN 4E091."562 Its logic was stated plainly: "this strategy aims to ensure that the model weights of the most advanced U.S. AI models are stored outside the United States only under stringent security conditions and the large clusters of advanced ICs necessary to train those models are built in destinations that pose comparatively low risks of diversion or misuse."563 It set numeric caps — "countries will be subject to a cumulative maximum installed base allocation of 790,000,000 TPP" from 2025 to 2027 — and per-company ladders calibrated, by the agency's own admission, to keep permitted clusters "approximately 12 months, or one generation, behind the cluster size BIS believes will be needed to train the most advanced dual-use AI models."564 It expressly left application programming interfaces untouched.565 On the thirteenth of May, 2025, the Department of Commerce announced its rescission, on the grounds that "[t]hese new requirements would have stifled American innovation and saddled companies with burdensome new regulatory requirements," and that the rule "would have undermined U.S. diplomatic relations with dozens of countries by downgrading them to second-tier status."566 A replacement rule was promised.566 I take the diplomatic objection seriously. But the rule was the only federal instrument that treated computation as a strategic physical quantity with a location and an owner — the same premise a compute register would have served domestically. Both instruments were withdrawn within four months of each other, and the reason offered in each case was the burden on the firms being measured.

XV.

I promised in the introduction that this book would judge power by a fixed standard rather than by whether I like the people holding it. The standard was this: no authority, public or private, elected or capitalized, may simultaneously control the means of force, of food and material provision, of communication and public truth, of transit and movement, of historical memory, of defining an emergency, and of predicting the future. Measured against that standard, on the disclosed evidence in this chapter: Five of the seven are held or effectively controlled by an overlapping set of four to six firms. A sixth — force — has been contractually promised to the state by those same firms rather than held independently of them, which is not the same thing as separation. I do not say this to be dramatic. The test was written before I looked, and the test does not care what I hoped to find.

XVI.

So here is what I would insist on. Not what I predict; what I would refuse to accept the absence of. Structural separation between infrastructure and applications. A firm that owns the compute on which rivals must train should not also compete with them in the application market that compute serves. This is the oldest rule in the American regulatory tradition, applied to railroads, pipelines, broadcast, and telephony, and abandoned exactly once — in the industry that most resembles all of them at once. The remedy need not be divestiture in the first instance. It can be a line of business restriction, or non-discriminatory access at published rates. It cannot be nothing. A compute register. Reinstate, by statute rather than executive order, a mandatory report of the location, ownership, and capacity of large-scale training clusters and of training runs above a stated threshold, with the threshold updated on a schedule rather than at discretion. The prior order's numbers were a reasonable start and its legal basis was fragile; a statute fixes both. The register does not have to be public in full. It has to

exist, and a regulator has to hold it. Public and academic compute at a scale that changes behavior. Not credits. Appropriated capacity, owned publicly, allocated by peer review, sufficient that an independent scientist can replicate a frontier result and an independent auditor can test a deployed system. Pass the authorizing statute that has now been reported out of committee without a dissenting vote. Interoperability and data portability as conditions of scale. Above a defined size, a provider must support the export of a customer's data, fine-tuned weights, and configurations in documented formats at cost, and must not degrade interfaces to competitors. Europe's rule is the template. Procurement rules with teeth against lock-in. The memorandum is correct and should be binding rather than advisory: publish the agreements, require exit terms and portability before award, prohibit nominal-price entry that is not paired with priced renewal terms, and forbid conditioning any discount on consolidation of spending. A government that cannot leave a vendor cannot regulate that vendor. Disclosure of cross-investments as a single map. Require every firm above a threshold to file, annually and in one place, its equity positions in suppliers and customers, its guarantees on their behalf, its committed purchases from them, and its revenue from them. Every element is already disclosed somewhere. Requiring it in one table would end the circularity's chief protection, which is that reconstructing it takes a reader three filings and a week.

Public-interest representation on any standards body whose output functions as law. If a private consortium's thresholds determine when a system is deployed, that consortium is exercising public authority and must carry public members with access to the underlying evidence. Alternatively, move the function to a public body. What is not acceptable is the current arrangement, in which the standard, the auditor, and the audited are the same six firms. And one procedural insistence. Every merger, partnership, licence, or talent transfer that transfers effective control of a frontier development team should be reportable, whatever its form. The agencies identified this loophole in 2021, published the finding, and did not close it. The form of a transaction should not determine whether the public gets to look at it.

XVII.

Everything in this chapter came from documents any citizen may read without permission, without a subscription, and without a lawyer. I want to be specific about how, because the habit is more useful than the conclusions. Start with the annual report on Form 10-K, filed with the Securities and Exchange Commission and free to the public. Four passages carry nearly all the weight. The concentration note tells you who the customers are, by percentage — this is where I found that two buyers account for thirty-six percent of a vendor's revenue.474 The commitments and contingencies note tells you what the company has promised to pay and to whom, and what it has guaranteed on someone else's behalf.478 The investments note tells you whose equity it holds, which is how you discover that a supplier owns pieces of its own customers.475 And the variable interest entity note tells you what it has built without putting on its balance sheet.500 A citizen who learns to find those four notes can reconstruct the structure of an industry in an afternoon. Second habit: read the filing against the press release. I have shown you two cases where the same firm said two different things about the same fact — a hundred billion dollars announced and then disclaimed as having "no assurance," and a stake described as roughly twenty-seven percent in a blog post and approximately twenty-five percent in an audited statement.480 487 Neither discrepancy is illegal. Both are informative. The document with legal liability attached to it is the document that tells you what the company believes. Third habit: search for absence. "Anthropic" appears zero times in Alphabet's annual and quarterly reports.488 "OpenAI" appears zero times in Oracle's annual report, in the same year that Oracle's remaining performance obligations rose to $638 billion.497 Absence is not proof of concealment; materiality is a judgment and counsel may have made it

honestly. But a relationship a firm will not name in a filing is a relationship you should ask about out loud. Fourth habit: court dockets are public and better than the coverage of them. Every quotation from the search, ad tech, and social platform cases in this chapter came from the opinion itself, retrieved free of charge. Reading the opinion rather than the article about the opinion is what let me tell you both that a court found a monopoly and that the same court declined the remedy that would have opened the market, stated its reasons, and put its own doubts on the record.519 521 No summary I read conveyed both halves. Fifth: at a hearing, ask the ownership questions rather than the technology questions. Who is the tenant of record, and who guarantees its lease? Does any equipment vendor hold equity in the tenant? Is the vendor's infrastructure the exclusive occupant of the site by contract? What happens to the guarantee if the tenant's funding does not materialize? Those questions are answerable from the record without any technical expertise at all. A four-gigawatt campus contractually closed to a vendor's competitors is a fact about your grid and your county, disclosed in a quarterly report, and available to be read into the minutes of a public meeting.481 I am not suggesting that citizens should have to do this. I am reporting that at present they must, because the government instrument that would have collected the same information in one place was created and then revoked, and no successor exists.533

XVIII.

Let me state the strongest objection to all of this, because it is a good one. Scale is not an accident of this industry; it is a property of it. Training a frontier model requires capital at a scale only a few institutions on earth can assemble. The firms in this chapter did not concentrate the market by conspiracy. They concentrated it by being the only entities willing to spend a hundred billion dollars on an uncertain return, and the returns have in fact produced tools that a great many people find genuinely valuable. Break them up and you may get five firms that can each afford a fifth of a training run, which is to say a world with no frontier models and the same problems. I accept the premise and reject the conclusion. The argument proves that scale may be necessary for capability. It does not touch the question of accountability, and it is deployed as though it did. Nothing about the capital requirements of training a model requires that the same firm own the cloud, the chips, the distribution channel, the standards body, and the equity in its own customers. Nothing about them requires that the compute register be abolished, that the training data be secret, that the acquisition be structured to avoid filing, or that the safety threshold be set by the party that profits from clearing it. Those are choices, made separately, each defensible on its own terms, and each pointing the same direction.

That is the pattern I most want citizens to see. No one decided to= concentrate this industry. A court declined a payment ban to protect= third parties. An administration revoked a reporting threshold to= reduce burden. A department rescinded an export rule to repair= alliances. A procurement agency accepted a nominal price to get= tools to public servants. A committee reported a bill and the floor= never scheduled it. Every one of those decisions is defensible in= isolation, and every one of them was made in isolation, and their sum= is a market structure no one chose and no one is accountable for. Cumulative decisions require cumulative review. I would require= every agency touching this industry — antitrust, procurement,= energy, security, export control — to publish, annually, an= assessment of the effect of its own decisions on market= concentration, and to state whether the aggregate direction is toward= or away from it. Not a new agency. A new sentence, required in= documents that already exist. I will end where the arithmetic ends. Scale can produce genuine= capability. It cannot produce accountability, and it must not be= permitted to substitute for consent. — D.M.D.

What Form Requires

1. Full public accounting of energy, water, land, emissions, waste, infrastructure, and community-health impacts.

2. Cumulative review alongside other projects sharing the same grid, watershed, coast, tax base, or labor market.

3. No permit without financial assurance sufficient to complete repair if the developer fails.

4. Avoidance of wetlands, aquifers, farmland, and vital habitat before mitigation is considered.

5. Public access to contracts, resource commitments, and the terms by which public costs are allocated.

6. Meaningful local standing to condition, challenge, or refuse projects that impose long-term risk.

A Practice for the Reader

Choose one major project proposed or operating near you. Find its permits, resource commitments, and the body that approved it. Ask: What else shares this grid, water, coast, tax base, or workforce? Who reviewed the total? What is promised? What is secured? Who pays if it fails? What brings the project back for review if it grows?

Coda: What We Protect Must Be Fairly Shared

The material costs of progress do not disappear because they have been placed out of sight. They return as higher bills, depleted aquifers, weakened infrastructure, degraded land, lost livelihoods, and communities asked to carry more than they receive.

Protection alone is not enough. A project may meet a narrow environmental standard while still shifting costs downward and moving wealth upward.

The next question is not only what a system takes. It is what it gives back: who shares in the value, who bears the risk, and whether technological progress makes ordinary life more secure or more precarious.

NEXT: PART FOUR — EXCHANGE

WORK, WAGES, AND FINANCIAL EQUALITY

The Distribution Question Is the Whole Question

TECHNOLOGY MUST ENLARGE HUMAN AGENCY, NOT REPLACE IT.

We live in an age of extraordinary capability. But capability is not wisdom. A system does not become just because it is efficient. A prediction does not become legitimate because it is accurate. A decision does not become democratic because it is automated.

Artificial intelligence is becoming part of the conditions under which people learn, work, receive care, seek housing, obtain benefits, communicate, and participate in public life. It is also built into the physical world: land, water, energy, infrastructure, labor, and communities asked to absorb costs in the name of progress.

This foundational manifesto establishes the six conditions of freedom: Perception, Will, Form, Exchange, Threshold, and What Citizens Do. Together, they form a civic blueprint for truth, democratic agency, ecological stewardship, economic fairness, human dignity, and accountable public action.

THE SIX CONDITIONS OF FREEDOM

PERCEPTION • WILL • FORM • EXCHANGE • THRESHOLD • WHAT CITIZENS DO

"A workable Earth is an Earth on which ordinary people can breathe clean air, drink safe water, find meaningful work, participate in decisions that affect their lives, and pass forward a world no less livable than the one they inherited."

