No Ceiling, No Denominator, No Exit:

Artificial Intelligence and the Architecture of American Energy Law

Travis Gilly · July 2026*

Abstract

For half a century the United States has run a federal program to make machines use less energy. It has never once run a program to make the country use less energy. The Energy Policy and Conservation Act regulates efficiency, which is a ratio, and the statute contains no authority to set a ceiling on aggregate consumption. That distinction was tolerable while the ratio and the total moved together. Artificial intelligence has pulled them apart. Google reports that the energy required to answer a median Gemini text prompt fell by a factor of 33 in twelve months, and reports in the same year that its electricity demand rose 37 percent, the largest annual increase in the company's history.

This Article argues that the gap is structural, and that it appears three times. First, coverage. The gate through which a new product enters the consumer program at 42 U.S.C. § 6292(b)(1)(B) is a fraction with households in the denominator, and the standard-setting gate at § 6295(l)(1) is metered the same way, down to an aggregate figure that is expressly an aggregate of household use. When the Department of Energy proposed in 2013 to cover computer servers, it satisfied the statutory test by computing the energy used by servers in houses. Industry objected that the agency lacked authority over the commercial units, the agency withdrew the determination, and no coverage determination for servers exists today. The same agency, reading a 1978 definition, built the regulatory category of computer room air conditioner out of statutory text that never named it, and now sets mandatory efficiency floors on the machine that cools the data center while setting nothing on the machines inside it. Second, disclosure. Google publishes watt-hours per prompt and withholds the prompt count it used as the divisor, and its chief scientist has said on the record that the company is not comfortable revealing that number. OpenAI has released no model-specific energy data since 2020. The aggregate is therefore not merely unregulated; it is uncomputable by anyone outside the firms. Third, allocation. Compute is financed by risk capital that shareholders bear, and the grid that serves it is financed by a rate base that households cannot leave.

The Article makes three contributions. It identifies the household denominator as the mechanism of EPCA's blindness rather than treating the absence of AI standards as agency inattention. It shows that the withheld denominator in corporate environmental disclosure reproduces the statute's own epistemic structure, so that the regime and the firm converge on publishing the quotient and suppressing the total. And it locates the resulting cost-allocation question in existing doctrine, where the term of art means what the Seventh Circuit says it means and not what the industry hears.

Keywords: Energy Policy and Conservation Act; energy conservation standards; coverage determination; data centers; artificial intelligence; cost causation; *Illinois Commerce Commission v. FERC*; *Loper Bright Enterprises v. Raimondo*; Jevons paradox; environmental disclosure.

---

*Travis Gilly is Founder and Executive Director of the Real Safety AI Foundation. ORCID: 0009-0007-2954-6313. Correspondence: t.gilly@ai-literacy-labs.org. The author discloses the use of speech-to-text software (Wispr Flow) and Anthropic's Claude as assistive technology for a diagnosed neurodevelopmental condition. All legal analysis, case synthesis, theoretical framing, doctrinal arguments, and conclusions are solely the author's own.

Contents

I Introduction 4
- A The Argument in Three Clauses 6
- B Three Literatures, and What This Article Adds to Each 7
  - 1 Rebound, and Why This Is Not a Rebound Argument 7
  - 2 Measurement, and the Narrower Claim Here 9
  - 3 Allocation, and the Work This Article Does Not Duplicate 10
  - 4 The Comparative Track 11
  - 5 What Sits in the Gap 12
- C Two Disclaimers, Stated Precisely 13
  - 1 This Is Not a Rate-Level Paper 13
  - 2 Cost Causation Is Not a Claim About What Caused Anyone's Bill 14
- D What the Argument Does Not Establish 15

II No Ceiling: The Statute Regulates a Ratio, and the Gate Is a Fraction 15
- A The Two Programs and the Asymmetry Between Them 15
- B The Coverage Gate Is a Fraction With Households in the Denominator 17
- C Every Metric in the Statute Attaches to a Unit 18
- D The 2013 Server Determination: Counting Servers in Houses 19
- E Regulating the Exhaust: The Department Built the Air Conditioner Category Itself 21
- F The Commercial Door Is Open 23
- G The July 2026 Proceeding, and Why the Ceiling Question Survives It 24

III No Denominator: The Disclosure Regime Reproduces the Statute's Blindness 26
- A Google Publishes the Quotient and Withholds the Divisor 27
- B OpenAI Withholds Both Terms 28
- C The Wrong Unit, and Congress Already Rejected It 30
- D The Convergence 31

IV No Exit: The Bet and the Ticket 31
- A The Fiber Analogy and Why It Is Not an Overbuilding Claim 32
- B Cost Causation, in Its Doctrinal Sense Only 33
- C What Follows, Stated Modestly 36
- D What Virginia Did in One Session 37
  - 1 The Gate Is Load 38
  - 2 The Tax Builds the Instrument 38
  - 3 The Rate Class Does the Allocation 39
  - 4 What Virginia Did Not Do 40

V What This Argument Does Not Show 40
- A The Racks Are Full 40
- B Training Burns Power With Nobody Watching, and Nobody Counts It Here 41
- C Most Data Center Load Is Not AI 41
- D Households Do Not Pay the Data Center's Electrons 41
- E The Doctrine Tolerates Crude Matching, and Presumes Network Benefit 42
- F The Department Is Not the Villain 42

VI Objections 43
A The Decoupling Objection 43
B The Plateau Objection 45
C The Overestimation Objection 46
D The Rebound Objection 47
VII Conclusion 47

I. Introduction

The Energy Policy and Conservation Act has been the federal government's principal instrument for reducing the energy consumption of manufactured goods since 1975.[1] Its program has four parts: testing, labeling, the establishment of federal energy conservation standards, and certification and enforcement.[2] By its own terms, and by the Department of Energy's account of its own authority, every one of those parts operates on a machine. None operates on a total.

This is the first thing to understand about the statute, and the last thing anyone says about it. EPCA measures efficiency. Efficiency is a ratio: work delivered over energy consumed, or its inverse. A ratio has no ceiling. It can be driven arbitrarily close to the thermodynamic limit while the quantity in the numerator grows without bound, and when that happens the statute registers success. The Secretary is required to pursue "the maximum improvement in energy efficiency that is technologically feasible and economically justified,"[3] and is separately forbidden from ever moving a standard backward.[4] What the Secretary is nowhere authorized to do is say that enough electricity has been used.

For fifty years this was an academic distinction. Appliance stocks grew slowly, per unit consumption fell, and the two effects roughly offset. Artificial intelligence has separated them with unusual clarity, and it has done so inside a single company's published numbers. Google reports that the energy consumed by the median Gemini Apps text prompt fell by a factor of

---

1. Energy Policy and Conservation Act, Pub. L. No. 94-163, 89 Stat. 871 (1975) (codified as amended at 42 U.S.C. §§ 6291-6317). The Department of Energy describes the statute as authorizing it "to regulate the energy efficiency of a number of consumer products and certain industrial equipment." Energy Conservation Program: Procedures, Interpretations, and Policies for Consideration of New or Revised Energy Conservation Standards and Test Procedures for Consumer Products and Certain Commercial/Industrial Equipment, 91 Fed. Reg. 42,034, 42,037 (proposed July 7, 2026) (to be codified at 10 C.F.R. pt. 430) [hereinafter 2026 Process Rule NOPR].

2. *Id.* at 42,037 (citing 42 U.S.C. §§ 6291, 6293, 6294, 6295, 6296, 6311, 6313, 6314, 6315, 6316).

3. 42 U.S.C. § 6295(o)(2)(A); *see also id.* § 6313(a)(6)(A)-(C) (parallel command for covered equipment); *id.* § 6316(a).

4. *Id.* § 6295(o)(1) ("The Secretary may not prescribe any amended standard which increases the maximum allowable energy use, or, in the case of showerheads, faucets, water closets, or urinals, water use, or decreases the minimum required energy efficiency, of a covered product."); *accord id.* § 6313(a)(6)(A)(iii)(I) (same rule for covered equipment).

33 over the twelve months ending in May 2025.[5] Google reports in the same window that its annual electricity demand rose 37 percent, which the company describes as its largest load growth in history.[6] Since 2019 the power consumed by Google's products and services has risen roughly 250 percent.[7] The company's own explanation is candid: its "AI infrastructure buildout is currently accelerating faster than the grid is decarbonizing."[8]

Both sets of numbers are true, and neither contradicts the other, because they measure different objects. The first is a ratio. The second is a total. EPCA can see the first. It has no instrument for the second, and this Article's claim is that the absence is architectural rather than accidental.

A reader who follows the energy literature will object here that the divergence is not established, because the leading work in the field holds that efficiency gains have absorbed demand growth in data centers and that forecasts to the contrary have a poor record.[9] That objection is serious and Part VI answers it at length. In short: the authors of the decoupling thesis attached expiration dates to it, they have since published that artificial intelligence has changed

---

5. Google, MEASURING THE ENVIRONMENTAL IMPACT OF DELIVERING AI AT GOOGLE SCALE 1 (2025), https://arxiv.org/abs/2508.15734 (last visited July 15, 2026) [hereinafter GOOGLE AI SCALE] (reporting "a 33x reduction in energy consumption and a 44x reduction in carbon footprint for the median Gemini Apps text prompt over one year"). The reduction is decomposed as roughly 23x from model and software improvements and 1.4x from improved hardware utilization. *See* Measuring the Environmental Impact of AI Inference, GOOGLE CLOUD BLOG (Aug. 21, 2025), https://cloud.google.com/blog/products/infrastructure/measuring-the-environmental-impact-of-ai-inference (last visited July 15, 2026) [hereinafter AI INFERENCE BLOG].

6. Read Google's 2026 Environmental Report, GOOGLE BLOG (June 30, 2026), https://blog.google/company-news/outreach-and-initiatives/sustainability/2026-environmental-report/ (last visited July 15, 2026) [hereinafter GOOGLE 2026 REPORT BLOG] (reporting Google's "largest load growth in history," a 37 percent annual increase in electricity demand, alongside a 2 percent year-over-year reduction in operational emissions).

7. Heather Clancy, *Google's AI Expansion Weighs Heavily on Climate Goals*, TRELLIS (June 30, 2026), https://trellis.net/article/googles-ai-expansion-weighs-heavily-on-climate-goals/ (last visited July 15, 2026) (reporting, from Google's 2026 environmental report, that "[s]ince 2019, the power consumed by its products and services has increased by 250 percent").

8. GOOGLE 2026 REPORT BLOG, *supra* note 6.

9. *See* Arman Shehabi et al., *Data Center Growth in the United States: Decoupling the Demand for Services from Electricity Use*, 13 ENV'T RSCH. LETTERS 124030 (2018), https://doi.org/10.1088/1748-9326/aaec9c [hereinafter Shehabi, *Decoupling*]; Eric Masanet et al., *Recalibrating Global Data Center Energy-Use Estimates*, 367 SCIENCE 984 (2020), https://doi.org/10.1126/science.aba3758 [hereinafter Masanet, *Recalibrating*]. Part VI takes the objection up in full.

the picture, and the argument here survives the thesis being correct, because whatever produced the decoupling it was not the federal efficiency program, which has never covered a server.

A. The Argument in Three Clauses

The title states the structure. Each clause is a Part.

No ceiling is a claim about coverage and about the metric. EPCA sets floors on efficiency and, in some standards, caps on the rated energy use of a single unit under an assumed duty cycle. It has never set, and contains no authority to set, a limit on how much energy a class of machines may consume in the aggregate. Part II establishes this from the text, and then shows something narrower and more damaging: the gate through which any new consumer product enters the program is itself a fraction whose denominator is households.[10] A machine that no household uses does not produce a small quotient. It produces no quotient. The statute cannot compute the quantity it requires itself to find.

No denominator is a claim about disclosure, and it is the Article's original contribution. Google's 0.24 watt-hours per prompt is a quotient.[11] Google computed it by dividing measured energy by a measured prompt count.[12] It publishes the quotient and withholds the divisor, and its chief scientist has explained why on the record.[13] Part III shows that this is the same epistemic move the statute makes, arrived at independently, and that the convergence is the finding.

No exit is a claim about allocation. Part IV distinguishes the compute layer, financed by risk capital whose holders bear their own losses, from the grid layer, financed by a rate base from

---

10. 42 U.S.C. § 6292(b)(2)(A) (defining "average annual per-household energy use" as "the estimated aggregate annual energy use (in kilowatt-hours or the Btu equivalent) of consumer products of such type which are used by households in the United States, divided by the number of such households which use products of such type").

11. GOOGLE AI SCALE, *supra* note 5, at 1 ("we find the median Gemini Apps text prompt consumes 0.24 Wh of energy").

12. *Id.* at 5 (describing the method of measuring "total user prompt count across Gemini Apps products for each LLM model" and dividing each energy component by that count over the same period).

13. Casey Crownhart, *Google's Still Not Giving Us the Full Picture on AI Energy Use*, MIT TECH. REV. (Aug. 28, 2025), `https://www.technologyreview.com/2025/08/28/1122685/ai-energy-use-gemini/` (last visited July 15, 2026) (quoting Jeff Dean: "We're not comfortable revealing that for various reasons," and reporting his explanation that the total "is an abstract measure that changes over time" and that "the company wants users to be thinking about the energy usage per prompt").

which residential customers cannot withdraw. The distinction is not a prediction that the buildout will fail. It survives demand arriving exactly on schedule, which is why it is the right one.

B. Three Literatures, and What This Article Adds to Each

Each clause of the title corresponds to an existing body of scholarship. The three bodies do not cite one another. They are housed, respectively, in machine learning conferences, in energy economics journals, and in energy law programs, and the organizing claim of this Article is that they are describing one architecture from three sides.

1. Rebound, and Why This Is Not a Rebound Argument

That efficiency gains need not reduce total consumption is among the oldest propositions in energy economics. Jevons stated it in 1865 about coal.[14] Its modern statement is Khazzoom's, and the subject of that statement was this statute's program. Writing in 1980, five years after EPCA, Khazzoom attacked the energy savings then being predicted from mandated appliance standards as "derived mechanically," on the assumption that a one percent improvement in efficiency yields a one percent drop in demand, and he cited the Department of Energy's own reports as the examples.[15] Brookes reached a parallel conclusion independently, and the pair of claims became the Khazzoom-Brookes postulate.[16] The magnitude has been contested ever

14. W. STANLEY JEVONS, THE COAL QUESTION: AN INQUIRY CONCERNING THE PROGRESS OF THE NATION, AND THE PROBABLE EXHAUSTION OF OUR COAL-MINES (1865).

15. J. Daniel Khazzoom, *Economic Implications of Mandated Efficiency in Standards for Household Appliances*, 1 ENERGY J. 21, 21 (1980) ("[T]he estimates of energy savings predicted to result from these mandated standards are derived mechanically. When mandated standards raise the appliance efficiency by 1 percent, demand is predicted to drop by 1 percent; when they raise efficiency by 2 percent, demand is predicted to drop by 2 percent; and so on."). Khazzoom identifies the Department of Energy's 1979 and 1980 demand reports, and a contemporaneous California Energy Commission staff report, as instances of the reasoning he criticizes. *Id.* The article uses "appliance" generically, "to cover household durables." *Id.*

16. *See* Harry D. Saunders, *The Khazzoom-Brookes Postulate and Neoclassical Growth*, 13 ENERGY J. 131 (1992) (formalizing the postulate within neoclassical growth theory). The exchange that produced the literature ran through the same journal. *See* J. Daniel Khazzoom, *Energy Saving Resulting from the Adoption of More Efficient Appliances*, 8 ENERGY J. 85 (1987); Amory B. Lovins, *Energy Saving Resulting from the Adoption of More Efficient Appliances: Another View*, 9 ENERGY J. 155 (1988); J. Daniel Khazzoom, *Energy Savings from More Efficient Appliances: A Rejoinder*, 10 ENERGY J. 157 (1989).

since and remains contested.[17] Luccioni, Strubell, and Crawford brought the frame to artificial intelligence, arguing that efficiency gains in AI "may paradoxically spur increased consumption" and that a focus on direct per-unit impacts "misrepresents AI's true climate footprint."[18]

The international political economy literature reached a parallel conclusion earlier and by a third route. Dauvergne's study of artificial intelligence in global supply chains concluded that efficiency gains in the middle of those chains are "rebounding into more production and consumption," and that "the micro-level gains from AI . . . are not going to add up to macro-level solutions" for the environmental consequences of the system they improve.[19]

Three fields have now said the same thing about the same gap, in 1980, in 2022, and in 2025, working from energy economics, from political economy, and from machine learning respectively. None of the three statements is a legal claim, and none of them reaches the mechanism this Article identifies. That is the shape of the opening.

The reader who knows that literature will expect what follows to be a rebound argument. It is not, and the distinction is the point.

Rebound is a claim about behavior and prices: consumption rises because efficiency makes the underlying service cheaper. It is an empirical claim, its magnitude is disputed, and it can be estimated. This Article's claim is prior to that one, and it is legal rather than economic. EPCA does not fail to bound aggregate consumption because consumers respond to falling costs. It fails because the gate through which a machine enters the program is a fraction with households in the denominator, and a server is in no household. Set every rebound elasticity in

---

17. *See* Steve Sorrell, *Jevons' Paradox Revisited: The Evidence for Backfire from Improved Energy Efficiency*, 37 ENERGY POL'Y 1456 (2009) (reviewing the arguments and concluding that the evidence for economy-wide backfire "is far from conclusive" while suggesting that rebound effects are larger than conventionally assumed); Lorna A. Greening, David L. Greene & Carmen Difiglio, *Energy Efficiency and Consumption: The Rebound Effect, a Survey*, 28 ENERGY POL'Y 389 (2000). This Article takes no position in that dispute.

18. Alexandra Sasha Luccioni, Emma Strubell & Kate Crawford, *From Efficiency Gains to Rebound Effects: The Problem of Jevons' Paradox in AI's Polarized Environmental Debate*, in PROCEEDINGS OF THE 2025 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 76 (2025), https://doi.org/10.1145/3715275.3732007.

19. Peter Dauvergne, *Is Artificial Intelligence Greening Global Supply Chains? Exposing the Political Economy of Environmental Costs*, 29 REV. INT'L POL. ECON. 696, 696 (2022), https://doi.org/10.1080/09692290.2020.1814381.

the literature to zero and nothing in Part II changes. That is the test of whether an argument is a rebound argument, and this one does not pass it, which is why it is worth making separately.[20]

The relationship runs the other way as well. If the rebound literature is right that efficiency alone will not bound the total, then the remedy it implies is a constraint on the total. Part II establishes that the principal federal instrument for machine energy consumption contains no such authority and no gate through which the load could pass to acquire one. The economics has been describing, for forty-five years, a problem that the statute is not built to answer, and the two have never been introduced.

2. Measurement, and the Narrower Claim Here

A second literature measures what AI consumes. It begins with the carbon accounting of model training,[21] extends to the recalibration of data center energy estimates,[22] and turns, more recently, to inference and deployment.[23] Its situation report is that transparency has fallen as the models have grown, and that the resulting vacuum is filled by inaccurate estimates which then propagate through media and scientific publication alike.[24]

The same conclusion has been reached from the corporate reporting side. Sætra argued in 2021 that the standards, metrics, and rating apparatus assembled under the heading of

---

20. Part VI.D returns to this, and to the objection that a legal article contributes nothing by entering an unresolved economic dispute.

21. *See* Emma Strubell, Ananya Ganesh & Andrew McCallum, *Energy and Policy Considerations for Deep Learning in NLP*, in PROCEEDINGS OF THE 57TH ANNUAL MEETING OF THE ASSOCIATION FOR COMPUTATIONAL LINGUISTICS 3645 (2019), `https://doi.org/10.18653/v1/P19-1355`.

22. *See* Eric Masanet et al., *Recalibrating Global Data Center Energy-Use Estimates*, 367 SCIENCE 984 (2020).

23. *See* Alexandra Sasha Luccioni, Yacine Jernite & Emma Strubell, *Power Hungry Processing: Watts Driving the Cost of AI Deployment?*, in PROCEEDINGS OF THE 2024 ACM CONFERENCE ON FAIRNESS, ACCOUNTABILITY, AND TRANSPARENCY 85 (2024), `https://doi.org/10.1145/3630106.3658542`.

24. Sasha Luccioni, Boris Gamazaychikov, Theo Alves da Costa & Emma Strubell, *Misinformation by Omission: The Need for More Environmental Transparency in AI* (June 18, 2025) (unpublished manuscript), `https://arxiv.org/abs/2506.15572` (observing that "transparency around the costs and impacts of these models has decreased" and tracing how "inaccurate or de-contextualized best-effort estimates" spread).

environment, social, and governance reporting does not capture the sustainability impacts of artificial intelligence, and he has since proposed a protocol intended to close the deficiency.[25]

That work is the closest prior treatment of the subject of Part III, and Part III does not repeat it. The claim there is not that data about AI's environmental impact is scarce, which is well established and not this Article's to establish. The claim is structural and specific: one particular quantity was measured, was used as a divisor, and was withheld while the quotient computed from it was published; the firm's chief scientist explained on the record why; and the resulting epistemic posture is the same one the statute adopts in § 6292(b)(2)(A), reached independently and for unrelated reasons. The scarcity literature calls for disclosure in general. Part III identifies a specific number, the use to which it was put, and the reason its absence reproduces the design of the record rather than marking a gap in it. That claim is available only to a reader holding the statute and the disclosure in view at once, which is why it has not been made.

3. Allocation, and the Work This Article Does Not Duplicate

The third literature is the newest and the closest to Part IV. Martin and Peskoe reviewed roughly fifty state regulatory proceedings and concluded that utilities can fund discounts to large computing customers by socializing the costs through prices charged to the public, operating through special contracts reviewed in "short and conclusory orders," through the gap between federal and state cost allocation, and through colocation.[26] Their paper is empirical, it is about rate design, and it is the state of the art on the question Part IV addresses.

---

25. Henrik Skaug Sætra, *A Framework for Evaluating and Disclosing the ESG Related Impacts of AI with the SDGs*, 13 SUSTAINABILITY 8503 (2021), `https://doi.org/10.3390/su13158503` (arguing that "none of these frameworks sufficiently capture the nature of the sustainability related impacts of AI"); *see also* Henrik Skaug Sætra, *The AI ESG Protocol: Evaluating and Disclosing the Environment, Social, and Governance Implications of Artificial Intelligence Capabilities, Assets, and Activities*, 31 SUSTAINABLE DEV. 1027 (2023), `https://doi.org/10.1002/sd.2438`.

26. ELIZA MARTIN & ARI PESKOE, EXTRACTING PROFITS FROM THE PUBLIC: HOW UTILITY RATEPAYERS ARE PAYING FOR BIG TECH'S POWER (Harv. Elec. L. Initiative, Mar. 2025), `https://eelp.law.harvard.edu/wp-content/uploads/2025/03/Harvard-ELI-Extracting-Profits-from-the-Public.pdf` (last visited July 15, 2026). Their treatment of the governing standard is at 7-9. *Id.* at 7 ("[C]ost causation' is a guiding principle in ratemaking that dictates consumer prices should align with the costs the utility incurs to provide service to that customer or group of similar ratepayers."); *id.* (explaining that by approving rates that roughly meet the standard, commissions "prevent 'undue discrimination' between utility ratepayers, a legal requirement that is typically specified in state law").

It is also silent on Parts II and III. Martin and Peskoe do not discuss EPCA, energy conservation standards, coverage determinations, or per-prompt measurement, and this Article does not duplicate their survey of the proceedings.[27] What Parts II and III supply is a reason the inquiry they document cannot be run to conclusion. They observe that cost shifting is difficult to detect and that the difficulty is partly one of transparency. Part IV.C names the doctrinal mechanism: the burden of demonstrating an imbalance between costs and benefits falls on the party challenging the allocation, the quantity that would discharge that burden is the product of a published ratio and a withheld count, and the count is held by the beneficiary. The observation that the calculation is hard becomes, on that account, a prediction rather than a complaint.

4. The Comparative Track

A parallel conversation is developing in Europe, principally around the environmental provisions of the Artificial Intelligence Act, and it has reached a structurally similar conclusion by a different route. Vogelaar argues that environmental protection in that instrument is "of limited extent due to ambiguous terminology, a focus on transparency, and the voluntary nature of the provisions that regulate sustainable action."[28]

The two regimes fail to reach the aggregate for opposite reasons, and the contrast is the useful part. The European instrument names environmental sustainability as an object of concern and then declines to make the relevant obligations binding. EPCA makes its obligations binding, enforceable, and ratcheted in one direction, and meters them per household, so that the machine at issue never reaches the point where the binding obligation would attach. One regime can see the total and will not act on it. The other would act and cannot see it. A reform proposal drawn

27. The paper contains no reference to the Energy Policy and Conservation Act, to efficiency standards, or to the per-prompt disclosures discussed in Part III. *See generally id.*

28. Irene Vogelaar, *Sustainable AI: The Energy-Accuracy Trade-Off in the Artificial Intelligence Act* (Mar. 20, 2025), in ENVIRONMENTS (Joshua C. Gellers & Henrik Skaug Sætra eds.), in OXFORD INTERSECTIONS: AI IN SOCIETY (Philipp Hacker ed., 2025), `https://doi.org/10.1093/9780198945215.003.0030`. On the broader framing of technology within earth system governance, see Joshua C. Gellers, *Emerging Technologies and Earth System Governance in the Anthropocene: Editorial*, 14 EARTH SYS. GOVERNANCE 100157 (2022), `https://doi.org/10.1016/j.esg.2022.100157` (editorial introducing a special issue on artificial intelligence and digitalization).

from either regime would be useless in the other, which is a reason to state the American problem in its own terms rather than importing a European diagnosis.

5. What Sits in the Gap

The claim of novelty in this Article is modest and it is a claim about connection rather than discovery. The rebound economists established that the ratio and the total can diverge. The measurement researchers established that the total is not being disclosed. The energy lawyers established that the costs are being allocated anyway. No one has observed that the statute's coverage gate is a household fraction, that the leading firm's disclosure is a quotient over a suppressed divisor, and that these are the same structure appearing at two levels of a single system. The three findings have been sitting in three rooms.

Sætra also states this Article's factual premise, in one sentence, as an aside. Surveying the tension between the technology's promise and its costs, he observes that "while AI might be used to make various systems more efficient, and thereby reducing their emissions, the data centers running these AI systems are increasing their emissions."[29] The ratio falls and the total rises, noted and passed over, in a book with seventeen development goals to cover. What follows here is the demonstration that the federal statute charged with machine energy consumption is built so that it cannot register the second half of that sentence.

There is also a name for the error the statute commits, and it was supplied from outside law. Sætra's account of artificial intelligence and sustainable development holds that "it is necessary to consider such systems in context, and not as some isolated and neutral tool," because an AI system is a part of a sociotechnical system whose institutions, structures, and economic and political arrangements are "all mutually dependent."[30] The alternative, which he takes from

29. HENRIK SKAUG SÆTRA, AI FOR THE SUSTAINABLE DEVELOPMENT GOALS 3 (2022).
30. *Id.*

Barley and rejects, is "[a]n isolationist approach to technology."[31] Attending to the context is what makes the indirect and second-order effects visible.

That is a methodological prescription, and it is addressed to analysts. Part II reports something the prescription does not anticipate. In the American energy conservation statute the isolationist theory is not one approach among the several an analyst might choose. It is the law. The metric attaches to a unit, the gate divides by households, the Secretary is directed to the maximum improvement in the efficiency of a covered product, and there is no provision under which the system the product belongs to becomes the object of regulation. A prescription to attend to context, delivered to an agency operating under EPCA, is a prescription the agency lacks authority to follow. That obstacle has gone unremarked because remarking on it requires holding an energy statute and a sustainability literature in view at the same time, and the two are not read by the same people.

C. Two Disclaimers, Stated Precisely

Because the argument crosses from coverage into allocation, two things must be fenced off before Part II, and the second matters more than the first.

1. This Is Not a Rate-Level Paper

The Article takes no position on whether electricity prices in the United States are too high, too low, or correct. It does not argue that any household's bill should be smaller. The allocation question in Part IV is a question about who bears the cost of load that the statute cannot see, and it would remain a live question at any price level, including a falling one. A reader who arrives at Part IV expecting an affordability brief will find that the argument does not depend on affordability and does not deliver one.

---

31. *Id.* at 4 (citing STEPHEN R. BARLEY, WORK AND TECHNOLOGICAL CHANGE (2020)). The contrast Sætra draws is between testing whether a system performs its assigned task and asking the prior questions: whether it is built on commonly available data, whether it is proprietary, and whether it "require[s] a level of computing power that is only available to a select few countries." *Id.*

2. Cost Causation Is Not a Claim About What Caused Anyone's Bill

This is the trap, and it is worth spending a subsection on because the entire industry response to allocation arguments runs through it.

"Cost causation" is a term of art in public utility ratemaking. It describes a principle governing the allocation of costs among rate classes: that those who are charged for a facility should receive benefits roughly commensurate with what they pay.[32] The word "causation" in that phrase does not carry its ordinary meaning. It does not assert that a particular facility caused a particular price increase. It asks whether a cost has been assigned to a class in proportion to the benefit that class receives.

The ordinary meaning is a different claim entirely: that data centers raised somebody's electric bill. This Article does not make that claim. It is an empirical claim about price formation, it is contested, it depends on jurisdiction and rate design and gas commodity costs and capacity market clearing, and it is unnecessary to the argument.

The two claims share a word and share nothing else. The rhetorical structure that follows is predictable enough to name in advance: an allocation argument is made, and an empirical answer is given. The respondent observes that rates are set by many factors, that data centers pay their tariffed rates, and that no causal attribution has been established. All of that can be true while the allocation question remains entirely unanswered, because the allocation question was never about attribution. Part IV therefore uses the term only in its doctrinal sense, and this subsection exists so that a reader who has not been told the difference does not supply the confusion without help.

---

32. *Ill. Commerce Comm'n v. FERC*, 576 F.3d 470, 476 (7th Cir. 2009) (Posner, J.) (holding that the Federal Energy Regulatory Commission "is not authorized to approve a pricing scheme that requires a group of utilities to pay for facilities from which its members derive no benefits, or benefits that are trivial in relation to the costs sought to be shifted to its members"). The pin is confirmed by the Seventh Circuit's own later quotation of the passage. *See Ill. Commerce Comm'n v. FERC*, 756 F.3d 556 (7th Cir. 2014) (quoting the passage in full and citing it to "576 F.3d at 476").

D. What the Argument Does Not Establish

Part V is devoted to this, and it is not a formality. The honest version of this Article concedes that built racks run at saturation and that graphics processor lead times are long; that training consumes power with no users attached, by design, and sits outside Google's per-prompt accounting by that paper's own terms;[33] that a majority of data center load is conventional cloud computing rather than AI inference; and that residential customers pay for shared infrastructure, capacity, and commodity rather than for a data center's electrons. Any version of this Article that suppressed those facts would be the version the industry would prefer to answer.

II. No Ceiling: The Statute Regulates a Ratio, and the Gate Is a Fraction

A. The Two Programs and the Asymmetry Between Them

EPCA runs two programs. Title III, Part B established the Energy Conservation Program for Consumer Products Other Than Automobiles, codified at 42 U.S.C. §§ 6291 to 6309.[34] Title III, Part C established the Energy Conservation Program for Certain Industrial Equipment, codified at §§ 6311 to 6317.[35] A "consumer product" is one distributed in commerce to any significant extent for personal use or consumption.[36] "Industrial equipment" is defined by exclusion from that category: it consumes energy, it is distributed to a significant extent for industrial or commercial use, and it is not a covered product.[37]

Each program has a list, and each list has a residual clause. Nineteen consumer products are named by Congress, from refrigerators through urinals, and the twentieth entry is "[a]ny other

33. GOOGLE AI SCALE, *supra* note 5, at 5 ("LLM training & data storage: This study specifically considers the inference and serving energy consumption of an AI prompt. . . . We leave the measurement of AI model training to future work.").

34. 2026 Process Rule NOPR, *supra* note 1, at 42,037 ("Title III, Part B of EPCA established the Energy Conservation Program for Consumer Products Other Than Automobiles."). Upon codification, Part B was redesignated as Part A. *Id.* at 42,037 n.2.

35. *Id.* at 42,037 (citing Pub. L. No. 95-619, tit. IV, § 441(a)); *id.* at 42,037 n.3 (noting redesignation as Part A-1 on codification).

36. 42 U.S.C. § 6291(1).

37. *Id.* § 6311(2)(A).

type of consumer product which the Secretary classifies as a covered product under subsection (b).”[38] Eleven equipment types are named, from electric motors through unfired hot water storage tanks, and the twelfth is “[a]ny other type of industrial equipment which the Secretary classifies as covered equipment under section 6312(b).”[39] Neither list names a computer. Neither names a server.

The residual clauses look symmetrical. They are not, and the asymmetry is the whole of Part II.

To classify a consumer product under § 6292(b)(1), the Secretary must determine two things: that classification “is necessary or appropriate to carry out the purposes of this chapter,” and that “average annual per-household energy use by products of such type is likely to exceed 100 kilowatt-hours (or its Btu equivalent) per year.”[40] To classify industrial equipment under § 6312(b), the Secretary must determine one thing: “The Secretary may, by rule, include a type of industrial equipment as covered equipment if he determines that to do so is necessary to carry out the purposes of this part.”[41] That is the entire provision. There is no quantitative gate of any kind. And the purpose it points to is stated in national terms: “to improve the efficiency of electric motors and pumps and certain other industrial equipment in order to conserve the energy resources of the Nation.”[42]

One door has a numeric lock metered at the household. The other has no lock at all and opens onto a national purpose. Which door a machine must pass through is determined by whether it is sold for personal use. That allocation was written for a world in which the energy and the benefit shared an address.

---

38. *Id.* § 6292(a)(1)-(20).
39. *Id.* § 6311(1)(A)-(L).
40. *Id.* § 6292(b)(1)(A)-(B). One kilowatt-hour is 3,412 British thermal units. *Id.* § 6292(b)(2)(B).
41. *Id.* § 6312(b).
42. *Id.* § 6312(a).

B. The Coverage Gate Is a Fraction With Households in the Denominator

The 100 kilowatt-hour threshold is ordinarily read as a de minimis rule, a filter to keep the Department from writing standards for the electric pencil sharpener. Read the definition and it is something else.

"[A]verage annual per-household energy use with respect to a type of product" means "the estimated aggregate annual energy use (in kilowatt-hours or the Btu equivalent) of consumer products of such type which are used by households in the United States, divided by the number of such households which use products of such type."[43] "Household" itself is left to the Secretary's rules.[44]

Both terms of the fraction are household-scoped. The numerator counts only the energy used by units "used by households in the United States." The denominator counts only households that use the product. A unit that no household operates contributes nothing to the numerator, and it brings no household to the denominator. It is not that such a unit produces a number below the threshold. It produces no number at all. The statute's test is a division problem, and for a machine that lives exclusively in a commercial building the division is undefined.

This is the mechanism, and it is worth being precise about why it matters more than the threshold's height. Congress could set the figure at one kilowatt-hour and the result would be identical, because the arithmetic never reaches a data center. The gate is not too high. It is pointed at the wrong meter.

The standard-setting provision confirms the reading and doubles it. Once a product is classified under the residual clause, the Secretary may prescribe a standard for it only on findings that "the average per household energy use within the United States by products of such type (or class) exceeded 150 kilowatt-hours (or its Btu equivalent) for any 12-month period," and that "the aggregate household energy use within the United States by products of such type (or class)

43. *Id.* § 6292(b)(2)(A).

44. *Id.* § 6292(b)(2)(C).

exceeded 4,200,000,000 kilowatt-hours (or its Btu equivalent) for any such 12-month period," along with findings on technological feasibility and the insufficiency of labeling.[45]

Note what the second finding is and is not. It is an aggregate, and Congress evidently wanted a national total in the picture. But it is "aggregate household energy use." Not aggregate energy use. Even where EPCA reaches for a total, it reaches for a total of household consumption. There are 4.2 terawatt-hours somewhere in the statute, and they are all measured in houses.

So the consumer path is locked twice, at coverage and at standard-setting, and both locks are the same lock.

C. Every Metric in the Statute Attaches to a Unit

A careful reader will object that EPCA does contemplate caps. The objection is right and the answer is narrow.

The statute repeatedly speaks of a "maximum allowable energy use" and of a "maximum quantity of energy use specified in such standard."[46] It directs the Secretary to determine "the maximum improvement in energy efficiency or maximum reduction in energy use that is technologically feasible for each type (or class) of covered products."[47] So EPCA can, and sometimes does, express a standard as a quantity rather than a ratio: kilowatt-hours per year for a refrigerator, rather than a coefficient of performance.

Every one of those quantities is per unit, and every one is computed under a test procedure that assumes a representative average use cycle.[48] A quantity per unit under an assumed duty cycle is a ratio wearing different units. It constrains what one machine may draw

---

45. *Id.* § 6295(l)(1)(A)-(D).

46. *Id.* § 6295(o)(1) ("maximum allowable energy use"); *id.* § 6295(r) (providing that a standard "may include any requirement which the Secretary determines is necessary to assure that each covered product to which such standard applies meets the required minimum level of energy efficiency or maximum quantity of energy use specified in such standard").

47. *Id.* § 6295(o)(2)(A).

48. *Id.* § 6293 (consumer product test procedures); *id.* § 6314 (covered equipment test procedures); 2026 Process Rule NOPR, *supra* note 1, at 42,037 (describing DOE's obligation to develop procedures measuring energy use "during a representative average use cycle or period of use").

while running as the test assumes it runs. It says nothing about how many machines exist, how hard they are run, or what they draw in aggregate. The rule against backsliding at § 6295(o)(1) protects the per unit figure and only the per unit figure.[49]

The seven factors that govern economic justification make the point from the other direction. The Secretary must weigh manufacturer and consumer economic impact, life-cycle operating cost savings against price increases, "[t]he total projected amount of energy, or as applicable, water, savings likely to result directly from the standard," any lessening of product utility, competitive effects as determined by the Attorney General, "[t]he need for national energy and water conservation," and other relevant factors.[50] The third factor is projected savings from a standard. The sixth is a need to be weighed. Neither is a quantity of energy the country may not exceed. There is a rebuttable presumption of economic justification when first-year savings exceed a third of the incremental purchase cost,[51] which is again arithmetic about one unit.

The conclusion is available on the face of the statute and needs no inference. EPCA is a per unit efficiency instrument. It has never been a consumption limit, and no provision of it can be made into one by an agency. That is not a criticism of the Department. It is a description of the tool Congress handed it.

D. The 2013 Server Determination: Counting Servers in Houses

The best evidence that the household denominator is doing real work is what happened when the Department tried to regulate servers, because the record shows the agency solving the division problem exactly as the statute required and producing an answer that no one could use.

On July 12, 2013, the Department published a notice of proposed determination that computer servers qualify as a covered product under Part A of Title III.[52] Part A is the consumer

---

49. 42 U.S.C. § 6295(o)(1); *accord id.* § 6313(a)(6)(A)(iii)(I).

50. *Id.* § 6295(o)(2)(B)(i)(I)-(VII); *see also id.* § 6313(a)(6)(B)(ii) (parallel seven-factor test for covered equipment).

51. *Id.* § 6295(o)(2)(B)(iii).

52. Energy Conservation Program for Consumer Products and Certain Commercial and Industrial Equipment: Proposed Determination of Computer Servers as a Covered Consumer Product, 78 Fed. Reg. 41,868 (proposed July 12, 2013) (Docket No. EERE-2013-BT-DET-0034, RIN 1904-AD03).

program. The Department was proposing to bring the machines in the data center through the household door, because in the residual-clause architecture there was no other door pointed at them that the Department chose to use.

Doing so required the Department to find that average annual per-household energy use by servers was likely to exceed 100 kilowatt-hours. In the companion proceeding the Department showed its work. It "calculated average household energy use for computer systems, in households that use the product, based on data from published literature and under the assumption that computer systems contain at least one computer or server, and possibly a UPS as well."[53] It found that "[t]he estimated annual energy consumption of individual servers ranges from approximately 1900-2100 kWh/yr for mass-produced volume servers," and that "[u]nder the assumption that households that use servers would have at most one of these types of servers, DOE estimated the average annual household energy use for households that use servers to be at least 1900 kWh/yr."[54] On that basis the Department "tentatively determine[d] that the average annual per-household energy use for computer systems is likely to exceed 100 kWh/yr, satisfying the provisions of 42 U.S.C. 6292(b)(1)."[55]

Read that finding against the fraction. The Department cleared a 100 kilowatt-hour bar by 1,800 kilowatt-hours. It did so by computing the energy drawn by servers sitting in houses, divided by the number of houses containing a server. Every rack in every data center in the United States contributed exactly nothing to the calculation, because no household used them and so they entered neither term. The statutory test was satisfied, and it was satisfied by a population of machines that is not the population anyone cares about. The gate opened, and it opened on the wrong side of the building.

It then closed. Cisco Systems, the Information Technology Industry Council, the Consumer Electronics Association, and the Telecommunications Industry Association

---

53. Energy Conservation Program: Proposed Determination of Computer and Battery Backup Systems as a Covered Consumer Product, 79 Fed. Reg. 11,345, 11,347 (proposed Feb. 28, 2014) (Docket No. EERE-2013-BT-DET-0035, RIN 1904-AD04) [hereinafter Computer Systems Proposal].

54. *Id.*

55. *Id.*

commented that the proposal "improperly attempts to combine a variety of consumer products, which DOE has authority to regulate, with entirely dissimilar commercial products that DOE does not currently have the authority to regulate."[56] On February 28, 2014, the Department withdrew the server determination.[57] Publishing the withdrawal the same day as the computer systems proposal, it accepted the commenters' premise in terms: "In light of these comments, DOE clarifies that the proposed scope of coverage for this rulemaking relates only to consumer products. Thus, this rule applies to those computer systems that are of a type which, to any significant extent, are distributed into commerce for personal use or consumption. See 42 U.S.C. 6291(1). These consumer products can be distinguished from commercial/industrial equipment, which are of a type not sold for consumer use. See 42 U.S.C. 6311(2)(A)."[58]

Nothing was finalized. The Department later described the whole episode in the past tense and in exploratory terms: "During 2014, DOE explored whether to regulate UPSs as 'computer systems.' "[59] Uninterruptible power supplies were eventually covered on a different theory. Servers were not covered at all. Twelve years later, EPCA's covered products list still names nineteen products and a residual clause,[60] its covered equipment list still names eleven types and a residual clause,[61] and no server appears on either.

E. Regulating the Exhaust: The Department Built the Air Conditioner Category Itself

The Department does have a mandatory efficiency standard for one machine in the data center. It is the air conditioner.

---

56. *Id.* at 11,346.

57. Energy Conservation Program: Proposed Determination of Computer Servers as a Covered Consumer Product, 79 Fed. Reg. 11,350, 11,350 (proposed Feb. 28, 2014) ("Department of Energy (DOE) withdraws for further consideration a proposed determination that computer servers (servers) qualify as a covered product under Part A of Title III of the Energy Policy and Conservation Act (EPCA), as amended"); *id.* ("In light of public comments received from interested parties addressing the nature and use of servers, DOE is withdrawing the Notice.").

58. Computer Systems Proposal, *supra* note 51, at 11,346.

59. Energy Conservation Program: Energy Conservation Standards for Uninterruptible Power Supplies, 85 Fed. Reg. 1447, 1454 (Jan. 10, 2020) (citing Computer Systems Proposal, 79 Fed. Reg. 11,345, and 79 Fed. Reg. 41,656 (July 17, 2014)).

60. 42 U.S.C. § 6292(a).

61. *Id.* § 6311(1).

Computer room air conditioners are covered equipment. They are regulated at 10 C.F.R. § 431.97, on the Net Sensible Coefficient of Performance metric, and every unit manufactured on or after May 28, 2024 must meet the applicable minimum.[62] The equipment classes are differentiated by heat rejection medium, net sensible cooling capacity, and airflow direction over the coil.[63] In the 2023 rule the Department found the ASHRAE Standard 90.1-2019 levels equivalent to existing federal standards for six equipment classes and more stringent for the remaining forty-six.[64]

The obvious observation is that federal law sets a floor on the machine that removes the heat and no floor on the machines that make it. That observation is available to anyone and it is not the point. The point is how the air conditioner got there.

Congress never named computer room air conditioners. EPCA defines "commercial package air conditioning and heating equipment" as "air-cooled, water-cooled, evaporatively-cooled, or water source (not including ground water source) electrically operated, unitary central air conditioners and central air conditioning heat pumps for commercial application," and sorts that category into small, large, and very large by cooling capacity.[65] That definition was written for buildings. The Department read it, decided a data center cooling unit fits inside it, and created the regulatory category by rule. Its own words are the cleanest statement of the move anywhere in the record:

> Pursuant to its authority under EPCA (42 U.S.C. 6313(a)(6)(A)) and in response to updates to ASHRAE Standard 90.1, DOE has established the category of CRAC,

---

62. 10 C.F.R. § 431.97(f)(2) (2026) ("Each computer room air conditioner manufactured on or after May 28, 2024, must meet the applicable minimum energy efficiency standard level(s) set forth in this paragraph (f)(2)."); *see* Energy Conservation Program: Energy Conservation Standards for Computer Room Air Conditioners, 88 Fed. Reg. 36,392, 36,392 (June 2, 2023) (adopting amended standards effective August 1, 2023 with compliance required May 28, 2024) [hereinafter CRAC Standards Rule].

63. Energy Conservation Standards for Computer Room Air Conditioners and Dedicated Outdoor Air Systems, 84 Fed. Reg. 48,006, 48,009 (proposed Sept. 11, 2019) (describing the class structure and noting that "DOE's current energy conservation standards for 30 equipment classes of CRACs are codified at 10 CFR 431.97").

64. CRAC Standards Rule, *supra* note 60, at 36,393.

65. 42 U.S.C. § 6311(8)(A)-(D); *see* 10 C.F.R. § 431.92 (2026) (parallel definitions).

> which meets the EPCA definition of "commercial package air conditioning and heating equipment," but which EPCA did not expressly identify.[66]

The Department did the same thing for dedicated outdoor air systems.[67] And the regulatory definition of the covered machine is written by reference to the uncovered one: a computer room air conditioner is equipment "[u]sed in computer rooms, data processing rooms, or other information technology cooling applications."[68]

So the position is this. The Department found a machine Congress had not named, held that a 1978 definition reached it, defined it by the room it serves, wrote a test procedure for it, and has maintained mandatory federal standards on it for over a decade. It did all of that on the commercial side of the statute, where the interpretive latitude is real and the numeric gate does not exist. It has never done any of it for the servers in the same room.

That is not a story about an agency that lacked authority. It is a story about an agency that exercised precisely this authority, on precisely this equipment category, in precisely this building, for the appliance in the corner.

F. The Commercial Door Is Open

Which returns to § 6312(b), and to what it would take.

A server is an article of equipment that in operation consumes energy, that is distributed in commerce to a significant extent for industrial or commercial use, and that is not a covered product.[69] It is industrial equipment on the face of the definition, and the 2014 commenters said

---

66. CRAC Standards Rule, *supra* note 60, at 36,396 (citing 10 C.F.R. §§ 431.92, 431.97).

67. 84 Fed. Reg. at 48,009 ("DOE has established additional categories of equipment that meet the EPCA definition of 'commercial package air conditioning and heating equipment,' but which EPCA did not expressly identify. These equipment categories include CRACs . . . and DOASes").

68. *Id.* The Department has since amended the definition to track how a manufacturer markets a model, consistent with AHRI Standard 1360-2022. CRAC Standards Rule, *supra* note 60, at 36,396.

69. 42 U.S.C. § 6311(2)(A) (defining "industrial equipment"); *see id.* § 6292(a) (covered products list, which does not include servers).

so in terms when they described commercial servers as products the Department did not then have authority to regulate as consumer goods.[70]

To bring it into the equipment program the Secretary would have to determine that doing so is "necessary to carry out the purposes of this part," and the purpose is conserving "the energy resources of the Nation."[71] There is no kilowatt-hour figure to clear, no household to divide by, and no aggregate threshold. The determination would then have to survive the ordinary constraints: a test procedure that is not unduly burdensome,[72] a standard that is technologically feasible and economically justified against the seven factors,[73] substantial evidence on review,[74] and the unavailability defense if interested persons show by a preponderance that performance characteristics substantially the same as those generally available would disappear.[75] Those are real constraints and this Article does not predict the outcome.

The claim is narrower and it is sufficient: the door has no lock on it, and in fifty years no Secretary has tried the handle for the machines that now drive the fastest load growth in the grid's history. The Department reached through that same door to invent a category for their air conditioner.

G. The July 2026 Proceeding, and Why the Ceiling Question Survives It

On July 7, 2026, the Department published a notice of proposed rulemaking revising the Process Rule, the appendix at 10 C.F.R. part 430, subpart C, appendix A that governs how the Department decides whether to regulate and at what level.[76] It proposes to make appendix A binding for

70. Computer Systems Proposal, *supra* note 51, at 11,347 (recording the position of Cisco, ITI, CEA, and TIA). The commenters' point was jurisdictional in the narrow sense: the Part A consumer program does not reach commercial equipment. *Id.* It was not a claim that the Part A-1 industrial program could not reach it.

71. 42 U.S.C. § 6312(a)-(b).

72. *Id.* § 6314(a)(2).

73. *Id.* § 6313(a)(6)(B)(ii)(I)-(VII).

74. *Id.* § 6306(b)(2); *see* 2026 Process Rule NOPR, *supra* note 1, at 42,044 (describing the standard).

75. 42 U.S.C. § 6313(a)(6)(B)(iii)(II)(aa).

76. 2026 Process Rule NOPR, *supra* note 1, at 42,034-35. The docket is EERE-2025-BT-STD-0001; the regulatory identification number is 1904-AF72. *Id.* at 42,034. Comments are due August 6, 2026, and the Department held a public webinar on July 15, 2026. *Id.*

certain actions, to add a definition of "significant energy savings," to reinstate a comparative "walk up" analysis and the clear and convincing evidence description, to include economic thresholds, and to revert much of the text to the 2020 version.[77] A companion request for information on the Department's analytic methods was published the same day, with comments due September 8, 2026.[78] The Department grounds the proposal in Executive Order 14154 and in the "best reading" of EPCA after *Loper Bright Enterprises v. Raimondo*, with support drawn from *Michigan v. EPA*.[79]

Two features of the proposal bear on this Article, and neither is the headline.

The first is an asymmetry in the proposed binding effect. The Department would make the procedures binding for "regulatory actions (actions that may increase ECS stringency relative to existing requirements)," while reserving "greater flexibility for other actions," on the reasoning that with deregulatory actions "hindsight may provide a wealth of information" and that the Department "does not believe that there is the same need for procedural protections for deregulatory actions as there is for regulatory actions."[80] Procedure would ratchet in one direction. Whatever else that is, it is a proposal about the rate at which floors may rise, and it confirms rather than disturbs the point that floors are the only object in view. The same posture is visible elsewhere in the program: in September 2025 the Department removed a set of amendments to 10 C.F.R. part 431 as required by the Congressional Review Act, reverting the affected sections.[81]

---

77. *Id.* at 42,034.

78. Energy Conservation Program: Review of DOE's Analytic Methods for Setting Energy Conservation Standards, 91 Fed. Reg. 41,578 (July 7, 2026).

79. 2026 Process Rule NOPR, *supra* note 1, at 42,037, 42,043 (citing *Loper Bright Enters. v. Raimondo*, 603 U.S. 369 (2024), and *Michigan v. EPA*, 576 U.S. 743 (2015)); Exec. Order No. 14,154, 90 Fed. Reg. 8353 (Jan. 29, 2025).

80. *Id.* at 42,043 (binding effect for regulatory actions); *id.* at 42,044 (reasoning as to deregulatory actions).

81. 90 Fed. Reg. 43,371 (Sept. 9, 2025); *see* Pub. L. No. 119-8; ELECTRONIC CODE OF FEDERAL REGULATIONS, 10 C.F.R. pt. 431 editorial note, `https://www.ecfr.gov/current/title-10/chapter-II/subchapter-D/part-431` (last visited July 15, 2026) (recording that, as of September 9, 2025, the affected sections reverted to the December 22, 2024 version).

The second is what the proposal is not. It does not touch coverage thresholds, because it cannot; § 6292(b)(1)(B) and § 6295(l)(1) are statutory. It does not create authority to cap aggregate consumption, because no rule could. It is a proceeding about how hard it should be to raise a ratio. The Secretary himself framed the stakes in the register the statute invites, on the subject of clothes dryers.[82]

The comment on the Secretary's example is that he has the physics right and it proves the opposite of what it is offered for. A dryer that finishes in one cycle is more efficient per unit of dry laundry than one that takes three passes, which is why Congress moved dryers to a metric denominated in finished work. The example concedes that the operation is the wrong unit and the outcome is the right one. Part III is about what happens when an industry reports the operation anyway.

III. No Denominator: The Disclosure Regime Reproduces the Statute's Blindness

Part II established that EPCA cannot see aggregate consumption because its instruments are ratios and its gate is a household fraction. This Part shows that the private disclosure regime, which is the source of every current public figure for AI energy consumption, has independently arrived at the same architecture. It publishes ratios. It withholds totals. And in the one case where the withholding has been explained on the record, the explanation is that the firm wants the public thinking about the ratio.

This is the Article's original contribution, and it requires no case.

82. *See* Stuart D. Kaplow, *Rollback of Federal Appliance Efficiency Standards Happening: What Businesses Need to Know*, `https://stuartkaplow.com/legal-library/environmental-law/rollback-of-federal-appliance-efficiency-standards-happening-what-businesses-need-to-know/` (last visited July 15, 2026) (quoting Secretary Chris Wright: "In America, you should be able to choose between a drying machine that takes multiple cycles to dry your clothes and one that does it on the first try").

A. Google Publishes the Quotient and Withholds the Divisor

The move described in this Section is not unique to the 2025 paper, and the firm has made it before in a way that shows why it matters. In 2022 a Google team including the chief scientist reassured the field that machine learning energy use "held steady at <15% of Google's total energy use for the past three years," and predicted a plateau followed by a decline.[83] A share is a ratio, its denominator went unreported, and the denominator has since roughly tripled.

In August 2025 Google published the first production measurement of per-prompt AI energy consumption by a major provider.[84] The paper deserves credit before it receives criticism. It is more comprehensive than what preceded it, and its central methodological insight cuts against the company's interest.

The measured figure is 0.24 watt-hours for the median Gemini Apps text prompt.[85] That figure includes active accelerator power, host system energy, idle machine capacity, and data center overhead.[86] Google also reports what the same prompt measures under the narrower convention the field had been using, which excludes processors, memory, idle machines, and overhead: 0.10 watt-hours.[87] The same prompt on the same hardware in the same month is either 0.10 or 0.24 depending on where the measurement boundary is drawn, a factor of 2.4

---

83. *See* David Patterson et al., *The Carbon Footprint of Machine Learning Training Will Plateau, Then Shrink*, 55 COMPUTER 18 (2022), https://doi.org/10.1109/MC.2022.3148714. Quotations from this work follow the authors' openly available manuscript, https://doi.org/10.36227/techrxiv.19139645.v3 (Mar. 3, 2022). Part VI.B takes up the objection this Article's argument draws from the prediction.

84. GOOGLE AI SCALE, *supra* note 5, at 1. The document is a technical report and had not been submitted for peer review at publication. *See* Georgia Butler, *Google: Median Gemini Prompt Uses 0.24 Watt Hours of Power and Consumes 0.26ml of Water*, DATA CTR. DYNAMICS (Aug. 22, 2025), https://www.datacenterdynamics.com/en/news/google-median-gemini-prompt-uses-024-watt-hours-of-power-and-consumes-026ml-of-water/ (last visited July 15, 2026).

85. GOOGLE AI SCALE, *supra* note 5, at 1.

86. *Id.* (describing an approach accounting for "the full stack of AI serving infrastructure … including active AI accelerator power, host system energy, idle machine capacity, and data center energy overhead"). By Google's decomposition, "the primary energy draw originates from the active AI Accelerator power (0.14 Wh, 58% of total) and the necessary host CPU & DRAM power (0.06 Wh, 25%)," with provisioned idle machines and data center overhead accounting for the balance. *Id.* at 7.

87. AI INFERENCE BLOG, *supra* note 5 ("When we apply this non-comprehensive methodology that only considers active TPU and GPU consumption, we estimate the median Gemini text prompt uses 0.10 Wh of energy …. This is an optimistic scenario at best and substantially underestimates the real operational footprint of AI."); *see* GOOGLE AI SCALE, *supra* note 5, at 7 tbl.1 (comparing the "Existing Approach" and the "Comprehensive Approach").

sitting entirely in a definitional choice. Google made the honest choice and said the other one understates. Hold that fact; Part III.C returns to it.

Now the arithmetic. Total energy consumed by Gemini equals energy per prompt multiplied by prompts. Google has both terms. It computed the first by dividing the second into measured energy: the method is to "measure total user prompt count across Gemini Apps products for each LLM model" and then to divide each energy component by that count over the same period.[88] The denominator is not an estimate the company lacks. It is an input the company measured, used, and then removed from the published result.

Asked for it, Google's chief scientist declined. Jeff Dean told MIT Technology Review, "We're not comfortable revealing that for various reasons."[89] He added that the total is an abstract measure that changes over time, and that the company wants users to be thinking about the energy usage per prompt.[90]

That is the regime in one sentence, from the person best positioned to state it. The quotient is published because the quotient is what the company wants considered. The divisor is withheld because the product of the two is the number the quotient exists to displace. Others have made the same observation in less specific terms.[91]

B. OpenAI Withholds Both Terms

If Google's posture is publishing the numerator, OpenAI's is publishing nothing that resolves.

---

88. GOOGLE AI SCALE, *supra* note 5, at 5.

89. Crownhart, *supra* note 13.

90. *Id.*

91. *See* Shaolei Ren, quoted in Butler, *supra* note 81 ("They're just hiding the critical information."); *Google Reveals How Much Energy a Single AI Prompt Uses*, ENERGYSAGE (Aug. 25, 2025), `https://www.energysage.com/news/google-ai-energy-use-electric-bill-impact/` (last visited July 15, 2026) (noting that Google "still hasn't revealed how many total Gemini queries it processes daily, making it impossible to calculate the AI's total energy footprint"); *AI Efficiency: A Gemini Reality Check*, GREENLY (Sept. 29, 2025), `https://greenly.earth/en-us/leaf-media/data-stories/ai-efficiency-a-gemini-reality-check-with-greenly` (last visited July 15, 2026) (identifying, among "[d]ata transparency gaps," that "[t]he study doesn't disclose key information such as total prompt volumes").

OpenAI has released no model-specific energy data since GPT-3 in 2020.[92] Its chief executive published a per-query figure of 0.34 watt-hours in a personal blog post, unattached to any named model and without documentation.[93]

Into that vacuum came an outside estimate. The University of Rhode Island's AI lab estimated that a medium-length GPT-5 response of roughly 1,000 tokens consumes an average of 18.35 watt-hours and up to 40, against 2.12 watt-hours for GPT-4.[94] The method was to multiply average response times by the power draw of assumed hardware configurations, using published data center efficiency metrics.[95] The spread between 18 and 40 for the same nominal task is the measurement doing what measurement does when the thing measured will not describe itself.

Then the estimate needed a denominator, and had none. To reach a total, the analysts multiplied by 2.5 billion daily requests, a figure taken from press reports rather than from OpenAI, producing an estimate of roughly 45 gigawatt-hours per day.[96]

Set the two postures beside each other. Google measured the denominator, used it as a divisor, published the quotient, and declined to release the divisor. OpenAI released neither, and

---

92. *See OpenAI's Lack of Transparency on GPT-5 Energy Needs Sparks Concern*, COMPUTING (2025), `https://www.computing.co.uk/news/2025/ai/openai-s-gpt-5-sparks-energy-use-debate` (last visited July 15, 2026) ("OpenAI has not released energy data for any of its models since GPT-3 in 2020, which had 175 billion parameters.").

93. *Id.* (reporting that Altman's June figures "were not tied to a specific model and lacked documentation"); *see also What's the Environmental Cost of an AI Text Prompt? Google Says It Has an Answer*, CBS NEWS (Aug. 21, 2025), `https://www.cbsnews.com/news/ai-environment-impact-study-energy-usage-google-gemini-prompt/` (last visited July 15, 2026) (reporting the 0.34 watt-hour figure).

94. *See* Jowi Morales, *ChatGPT 5 Power Consumption Could Be as Much as Eight Times Higher Than GPT 4*, TOM'S HARDWARE (Aug. 14, 2025), `https://www.tomshardware.com/tech-industry/artificial-intelligence/chatgpt-5-power-consumption-could-be-as-much-as-eight-times-higher-than-gpt-4-` (last visited July 15, 2026) (reporting the URI AI lab estimates of 18.35 watt-hours average and up to 40 watt-hours, against 2.12 for GPT-4, 25.35 for o3, and 20.90 for DeepSeek R1).

95. *See OpenAI Flagship: Performance, Cost, Electricity*, DIGITIMES (Aug. 15, 2025), `https://www.digitimes.com/news/a20250815PD238/openai-flagship-performance-cost-electricity.html` (last visited July 15, 2026) (describing the URI team's derivation from response times, assumed graphics processor configurations, and power usage effectiveness factors, and noting that "OpenAI has not disclosed deployment specifics"). Researchers caution that extended reasoning modes can multiply the figure five to tenfold. *Id.* (quoting Shaolei Ren).

96. Morales, *supra* note 92 ("The University of Rhode Island's AI lab estimates that GPT-5 averages just over 18 Wh per query, so putting all of ChatGPT's reported 2.5 billion requests a day through the model could see energy usage as high as 45 GWh."). Other renderings of the same arithmetic express the result as the daily electricity use of about 1.5 million United States homes. *See* COMPUTING, *supra* note 90.

outside researchers reconstructed the numerator from hardware assumptions and borrowed the denominator from journalism. In both cases the aggregate is unavailable to the public, and in both cases it is available inside the firm. The information asymmetry is total and it is not incidental to the disclosure; it is the shape of the disclosure.

C. The Wrong Unit, and Congress Already Rejected It

The deeper problem with watt-hours per prompt is not that the denominator is missing. It is that a prompt is the wrong thing to divide by, and American energy law worked this out decades ago on a machine in the laundry room.

A prompt is an operation. It is not an outcome. The measure cannot see a retry, a failed answer, a rephrasing, a session in which the user asked six times to get one usable result. Every one of those is more prompts, so every one of them lowers the average while raising the total. A metric denominated in operations improves when the operation is performed badly.

Congress confronted exactly this and chose the other side. The Secretary is directed to develop procedures that measure energy use "during a representative average use cycle or period of use,"[97] and for clothes dryers the applicable metric is denominated in finished work rather than in machine operations, which is why the Secretary's own defense of the current proposal turns on a dryer "that takes multiple cycles" against one that finishes "on the first try."[98] A dryer rated per tumble would post an excellent number while the laundry stayed wet. Everyone in the program understands this. It is the reason the statute speaks of a use cycle rather than a motor revolution.

The AI industry reports per tumble.

And it does so while conceding, in the same document, that the boundary of the measurement is a choice. Recall the 0.10 against the 0.24.[99] Google's paper is admirable precisely because it says the narrow number understates and publishes the wide one anyway. But the concession establishes the general proposition that a per-operation figure is a construct

97. 2026 Process Rule NOPR, *supra* note 1, at 42,037 (describing the requirement of 42 U.S.C. §§ 6293, 6314).

98. Kaplow, *supra* note 80.

99. *See supra* notes 84-85 and accompanying text.

whose value depends on decisions the reporting party makes and the reader cannot audit. When the same party also declines to publish the count, there is nothing left for an outsider to check. A quotient whose numerator is boundary-dependent and whose denominator is confidential is not a disclosure. It is a number.

D. The Convergence

Put Parts II and III together, because the pairing is the finding and neither half produces it alone.

EPCA asks: what does one machine draw per household per year? It publishes floors on that quantity, and it has no line anywhere for the total. Google asks: what does one prompt draw? It publishes that quantity, and it withholds the count that would yield the total.

These are the same instrument. One was drafted in 1975 by a Congress responding to an oil embargo and codified through amendments across five decades. The other was designed in 2025 by an engineering organization measuring its own product. They were not copied from each other. They converged, because a ratio is what an institution reports when it wants credit for improvement without accountability for scale, and because efficiency improvements are the only kind of environmental progress that is compatible with unbounded growth.

The convergence explains the otherwise puzzling fact that the statute and the firm produce the same blind spot from opposite directions. The statute cannot see the data center because its denominator is households. The firm's disclosure cannot see the data center's total because its denominator is confidential. In both cases the ratio is published, the aggregate is not computed, and the institution is able to report that the number went down in a year when the number went up.

IV. No Exit: The Bet and the Ticket

Everything to this point is an account of what cannot be seen. This Part asks the legal question that the invisibility creates. The chain is short. The statute cannot see the load. The load is

therefore not constrained by the statute. The grid absorbs it. Once the grid absorbs it, someone pays for the absorption, and the allocation of that cost is governed by doctrine.

A. The Fiber Analogy and Why It Is Not an Overbuilding Claim

The episode invoked here is invoked in the energy literature for the opposite purpose, and the difference is worth stating at the outset. Koomey and Masanet treat the 1990s network buildout as the cautionary example of extrapolation error, an instance in which projecting from traffic "doubling every 100 days" produced "vast overinvestment in network capacity."[100] This Article draws nothing from the forecast. It draws on how the buildout was financed and who absorbed the loss.

The reference case is telecommunications in the 1990s, and the analytic frame for it is Carlota Perez's account of how technological revolutions are financed.[101] Carriers laid on the order of 80 million miles of fiber in the United States, financed substantially by debt, and most of it sat unlit for years after the crash.[102]

---

100. *See* Jonathan Koomey & Eric Masanet, *Does Not Compute: Avoiding Pitfalls Assessing the Internet's Energy and Carbon Impacts*, 5 JOULE 1625, 1625 (2021), https://doi.org/10.1016/j.joule.2021.05.007. Part VI.C addresses the critique, which this Article accepts and does not trigger.

101. CARLOTA PEREZ, TECHNOLOGICAL REVOLUTIONS AND FINANCIAL CAPITAL: THE DYNAMICS OF BUBBLES AND GOLDEN AGES (2002). The work is cited for its thesis as a whole. Perez divides each technological surge into an installation period, comprising irruption and frenzy, and a deployment period, comprising synergy and maturity, separated by a turning point that ordinarily takes the form of a financial crisis. The proposition relevant here is that the speculative frenzy is what installs the infrastructure the later deployment period uses, so that the capital destroyed in the crash and the capacity bequeathed to the successor are two descriptions of one event.

102. The magnitudes are contested and no authoritative accounting exists. Published estimates of the share left dark range from roughly 85 to 95 percent, and estimates of capital deployed range from roughly $1 trillion to $2 trillion. *Compare* Dan Runkevicius, *Is the AI Boom Headed for Its "Dark Fiber" Moment?*, FORBES (Mar. 24, 2025), https://www.forbes.com/sites/danrunkevicius/2025/03/24/is-the-ai-boom-headed-for-its-dark-fiber-moment/ (last visited July 15, 2026) ("By various estimates, 85% to 95% of fiber laid in the '90s remained unused after the bubble burst."), *with* Jim Hayes, *The Perils of Irrational Exuberance: The 25th Anniversary of the Dot-Com Boom*, ISE MAG. (July 24, 2025), https://www.isemag.com/professional-development-leadership/article/55293922/the-perils-of-irrational-exuberancethe-25th-anniversary-of-the-dot-com-boom (last visited July 15, 2026) (estimating as much as $2 trillion spent building 80 to 90 million miles of fiber network from 1995 to 2000, and 95 percent dark in 2001). An industry publication reports the underlying difficulty directly: much of the capacity "remains unlit, though nobody has a clear idea of how much." Logan Kugler, *Dark Fiber Is Lighting Up*, COMM. ACM (2014), https://cacm.acm.org/news/dark-fiber-is-lighting-p/ (last visited July 15, 2026). The comparison drawn in this Part is to the financial structure of the episode and not to its magnitudes, and the argument does not turn on which estimate is right. *See infra* Part IV.A.

The usual lesson drawn is that the industry built too much. That lesson is unavailable here and this Article does not draw it. Built racks are running at saturation, and lead times for accelerators run most of a year or more.[103] There is no evident dark capacity. Demand is arriving.

The structural feature worth carrying forward is different, and it survives demand arriving exactly on schedule. In the fiber episode, the party that made the bet and the party that held the ticket were the same party. Carriers borrowed, carriers built, and when the capacity went unsold the carriers and their investors and the pension funds behind them absorbed the loss. The bet and the ticket were in the same hand. That is what a functioning capital market looks like, and the misallocation was expensive without being a legal problem.

The AI buildout has two layers, and the hand changes between them.

The compute layer is risk capital. Hyperscalers are financing accelerators and buildings with shareholder money. If the demand fails to appear, or appears at prices that do not clear, shareholders take the loss. That is the arrangement they agreed to and it raises no question this Article needs to answer. It is, in the sense that matters here, none of the law's business.

The grid layer is a rate base. Generation, transmission, and the associated infrastructure are capitalized and recovered through regulated rates from a customer body that includes every residential customer in the service territory. A household cannot decline. There is no competitive substitute for the distribution wire, and the obligation to serve runs in both directions.

The bet and the ticket have separated. And the asymmetry with fiber is sharper than the separation alone suggests: dark fiber waited in the ground at approximately zero marginal cost, which is why the capacity could sit for a decade and then be lit profitably. Stranded generation does not wait at zero cost. A combustion turbine built for a load that does not materialize accrues carrying costs, and those costs sit in a rate base.

B. Cost Causation, in Its Doctrinal Sense Only

Here the terminology discipline of Part I.C.2 has to hold.

---

103. *See infra* Part V (setting out the concessions in full).

The cost causation principle in federal ratemaking asks whether costs assigned to a class are roughly commensurate with the benefits that class receives. The Seventh Circuit's statement is the standard one: the Commission "is not authorized to approve a pricing scheme that requires a group of utilities to pay for facilities from which its members derive no benefits, or benefits that are trivial in relation to the costs sought to be shifted to its members."[104] The principle has been applied and refined since.[105]

Two features of the doctrine matter for the argument, and the second is where the industry's move lives.

First, the inquiry is comparative and it is about allocation. It does not ask whether a facility was worth building. It asks whether the people billed for it are the people benefiting from it, in a rough proportion. A facility can be enormously valuable and still be misallocated, and the doctrine is indifferent to how good the thing is.

Second, "benefit" in this doctrine is a term with a technical referent. In transmission cost allocation it concerns load flows, reliability, and the physical and economic effects on an identified set of customers. It does not mean benefit in the sense of contribution to human welfare, scientific progress, or economic dynamism.

That distinction is the load the industry's rhetoric is built to carry across. The argument runs: artificial intelligence produces enormous collective benefit, therefore the costs of the infrastructure serving it are properly socialized, because collective benefit is what justifies collective cost recovery. Every step of that is a word game. The premise uses "benefit" in the

---

104. *Ill. Commerce Comm'n*, 576 F.3d at 476.

105. *See Old Dominion Elec. Coop. v. FERC*, 898 F.3d 1254, 1256 (D.C. Cir. 2018) (framing the principle as one ensuring that "burden is matched with benefit," so that the Commission "generally may not single out a party for the full cost of a project, or even most of it, when the benefits of the project are diffuse" (quoting *BNP Paribas Energy Trading GP v. FERC*, 743 F.3d 264, 268 (D.C. Cir. 2014))); *S.C. Pub. Serv. Auth. v. FERC*, 762 F.3d 41, 53 (D.C. Cir. 2014) (per curiam) (stating the Order No. 1000 formulation, that transmission costs be allocated to regional beneficiaries "in a manner that is at least roughly commensurate with estimated benefits"); *Ill. Commerce Comm'n v. FERC*, 721 F.3d 764 (7th Cir. 2013) (affirming approval of a regional cost allocation where precise quantification was impossible, on the reasoning that if crude matching "is all that is possible, it will have to suffice"). Compliance with Order No. 1000 does not by itself satisfy the principle, which the D.C. Circuit describes as "a pre-existing, more general rule that, in order to ensure just and reasonable rates, FERC must make some reasonable effort to match costs to benefits." *Old Dominion*, 898 F.3d at 1256.

mission-statement sense. The conclusion needs it in the load-flow sense. The doctrine supplies the conclusion only on the technical reading, and on the technical reading the premise is not evidence of anything.

The doctrine has already answered the move, in the case that states it. The Seventh Circuit endorsed the Commission's own conclusion that "a claim of generalized system benefits is not enough" to justify shifting costs onto a class where the proponent "point[s] to no evidence in the record that seeks to quantify this benefit, or even shows that such a benefit has occurred," so that the alleged benefits are "too speculative and unsupported to be taken into account."[106] Generalized system benefits are the entire content of the claim that artificial intelligence benefits everyone. Stated in a rate proceeding, in those words, it is the argument the Commission rejected and the court approved rejecting.

The claim also has a literature, and that literature is not friendly to it. Dauvergne found that the sustainability gains firms attribute to artificial intelligence are exaggerated by the metrics and rhetoric of corporate social responsibility, and that presenting the technology as a force for sustainability operates to legitimize business as usual and to obscure the case for state regulation.[107] That is a finding in political economy, not a fact in any docket, and it proves nothing about any allocation. It is a reason to expect the generalized-benefits claim to arrive, and to arrive in precisely the register the doctrine has already held insufficient.

Cost causation survives artificial intelligence being wonderful. If the technology cures disease and doubles productivity and does everything its proponents say, the allocation question is untouched, because the question was never whether the thing is good. It was who is on the invoice.

---

106. *Ill. Commerce Comm'n*, 576 F.3d at 476 (quoting *Transcontinental Gas Pipe Line Corp.*, 112 F.E.R.C. ¶ 61,170, at 61,924-25 (2005)).

107. Dauvergne, *supra* note 19, at 696.

C. What Follows, Stated Modestly

This Article does not propose a rate design, and it does not assert that any existing allocation violates the principle. The commissions are the fact-finders and the records are theirs.

What follows is smaller and it is the point of putting the three Parts in one Article. When a load is invisible to the statute that regulates energy use, unconstrained in aggregate by any instrument, and unquantifiable from public disclosure, the allocation question cannot be answered on the merits by anyone outside the firms, because the inputs are held by the firms. The doctrine asks for a comparison of costs to benefits across classes. Performing that comparison requires knowing the magnitude of the load, its growth path, its utilization, and its relationship to the infrastructure being capitalized. The hyperscalers have those numbers. Google has the prompt count. Neither publishes.

There is a further step, and it is the one that makes the three Parts a single argument rather than three observations. The doctrine places the burden on the challenger. Reviewing an allocation it went on to uphold, the Seventh Circuit disposed of the objection in a sentence: the challenger "can't counter FERC without presenting evidence of imbalance of costs and benefits, which it hasn't done."[108] That is a sensible rule when the evidence exists somewhere a challenger can reach.

Evidence of imbalance between costs and benefits requires the magnitude of the load. The magnitude of the load is the product of a ratio and a count. The ratio is published. The count is confidential, and the firm that holds it has said it is not comfortable revealing it.[109] So the doctrine assigns the burden of proof to the party that the disclosure regime disables, and assigns it with respect to a quantity that only the beneficiary of the allocation can compute. That is not a

---

108. *Ill. Commerce Comm'n v. FERC*, 721 F.3d 764 (7th Cir. 2013). The court also held that crude matching suffices where precision is unavailable, reasoning that "if crude is all that is possible, it will have to suffice." *Id.* The citator carries a caution signal on this decision; it is cited here for the burden and sufficiency propositions, which later panels have not disturbed. *See Long Island Power Auth. v. FERC*, 27 F.4th 705 (D.C. Cir. 2022).

109. Crownhart, *supra* note 13.

claim that any allocation is wrong. It is a claim that the mechanism the doctrine relies on to find out is, for this category of load, inoperative.

The missing denominator of Part III is therefore not only an environmental reporting problem. It is a defect in the evidentiary predicate of a legal doctrine that is currently being applied, in dozens of proceedings, on records the regulated parties construct.

D. What Virginia Did in One Session

The preceding Sections describe an allocation problem that federal law is not positioned to solve, and a reader is entitled to ask whether the problem is soluble at all. It is. A state solved most of it while this Article was being written, and the manner of the solution identifies the federal defect with more precision than any argument advanced here.

On June 30, 2026, the Governor of Virginia signed the Commonwealth's biennial budget, which created a Data Center Electricity Consumption Tax effective the following day.[110] The measure imposes a levy of $0.011 per kilowatt-hour on all electricity consumed at each Virginia data center each month, applies to utility-supplied and self-generated power alike, caps general-fund revenue at $600 million annually with the excess refunded pro rata, and sunsets on June 30, 2028.[111] Virginia is the largest data center market in the United States, and no state had previously charged a tax on the electricity a data center consumes.

Three features of the enactment bear on the argument of this Article, and the first is the one that matters most.

---

110. H.B. 30, 2026 Gen. Assemb., Spec. Sess. I, item 3-5.24 (Va. 2026) (enacted) [hereinafter Item 3-5.24], `https://budget.lis.virginia.gov/bill/2026/2/HB30/Chapter/`. The 2026 regular session adjourned without a budget; the enacted bill is the Special Session I instance of H.B. 30, and the regular-session instance of the same number was continued to the special session and never became law. *See* Bradley J. Nowak et al., *Virginia Budget Creates New Electricity Consumption Tax for Data Centers*, Williams Mullen (July 2026), `https://www.williamsmullen.com/insights/news/legal-news/virginia-budget-creates-new-electricity-consumption-tax-data-centers`.

111. Item 3-5.24, *supra* note 110. Utilities collect and remit monthly; operators who self-supply remit quarterly, with first payments due in September 2026. Reporting that the rate is $0.11 per kilowatt-hour is in error by a factor of ten. *Compare* Giorno et al., *infra* note 115, *with* Nowak et al., *supra* note 110 ($0.011).

1. The Gate Is Load

The tax reaches a facility "primarily engaged in the storage, management, and processing of digital data" that supports "at least one megawatt of electrical capacity."[112]

That is a coverage determination, and it took the General Assembly one budget cycle to write. Its jurisdictional threshold is electrical capacity. A machine is inside the class if it draws a megawatt, and the drawing of the megawatt is the entire inquiry.

Set that beside 42 U.S.C. § 6292(b)(1)(B), which conditions coverage on a finding that the type of product "is not consumed at [certain] levels" and, at subsection (b)(2)(A), instructs the Secretary to compute average annual per-household energy use by dividing aggregate consumption by the number of households in which the product is found.[113] The two provisions ask different questions about the same machine. Virginia asks how much power it draws. EPCA asks how many households own it. A server draws megawatts and sits in no household, so it clears one gate on the first question and fails the other before the first question is reached. The federal statute's silence is therefore not evidence that the object is hard to regulate. It is evidence about the gate.

2. The Tax Builds the Instrument

Part III described a firm that published a rate per prompt and withheld the prompt count.[114] The measurement that would close that gap is not technically exotic. It is a meter, and Virginia has now required one, not because it set out to solve a disclosure problem but because a per-kilowatt-hour levy cannot be assessed without counting kilowatt-hours.

The legislator who described the mechanism said so directly. Because the statute caps what the Commonwealth keeps rather than what it collects, the metering continues past the point

---

112. Item 3-5.24, *supra* note 110; *see* Nowak et al., *supra* note 110 (reciting the enacted definition). The definition encompasses computer systems, networking equipment, electrical infrastructure, cooling systems, environmental controls, and fire suppression systems, and excludes facilities whose primary function is internet access service or communications service under VA. CODE ANN. § 58.1-647.

113. 42 U.S.C. § 6292(b)(1)(B), (b)(2)(A); *see supra* Part II.B.

114. *See supra* Part III.B.

where the revenue stops, and the consequence, in his words, is that "we're actually going to know how much electricity these folks are using."[115]

The Governor's own recommendation on the enrolled bill drove the point further. Amendment 11 to Item 3-5.24 requires that a supplier "list the applicable tax as a separate line item on the data center operator's billing invoice."[116] The consumption is therefore not only metered. It is itemized, monthly, on a document the operator receives.

A revenue measure produced, as a byproduct, the disclosure that a decade of voluntary sustainability reporting did not. The number exists. It was always going to exist the moment somebody had a reason to bill for it.

3. The Rate Class Does the Allocation

On the allocation question the Commonwealth moved twice. In November 2025 the State Corporation Commission approved a new rate class for customers drawing more than 25 megawatts, effective January 2027, under which such customers pay minimums of 85 percent of contracted transmission and distribution demand and 60 percent of generation demand.[117] The legislature then enacted S.B. 253, which authorizes the Commission to assign the electricity costs driven by large users to those users rather than to residential and small business ratepayers.[118] Cost causation, which Part IV.B describes as a doctrine the federal regime articulates and then declines to operationalize against the party holding the data, is in Virginia an instruction to a commission with a number attached.

---

115. Taylor Giorno, Zara Norman & Samuel Larreal, *'The Battle Line Has Been Drawn' Around Virginia's Data Centers*, NOTUS (July 1, 2026), `https://www.notus.org/metro/northern-virginia-data-centers-politics` (quoting Del. John McAuliff) [hereinafter Giorno et al.].

116. 2026 Special Session I, Budget Amendments, H.B. 30 (Governor's Recommendations) amend. 11 (June 26, 2026), `https://budget.lis.virginia.gov/amendment/2026/2/HB30/Enrolled/GR`. The amendment is described as a technical adjustment clarifying billing.

117. *See Virginia's New Data Center Electricity Rate Class*, Am. Action F. (Apr. 22, 2026), `https://www.americanactionforum.org/insight/virginias-new-data-center-electricity-rate-class/`.

118. S.B. 253, 2026 Gen. Assemb., Reg. Sess. (Va. 2026). *See Virginia Leads the Way on Energy Affordability in 2026*, Nat'l Caucus of Env't Legislators (Apr. 2, 2026), `https://www.ncelenviro.org/articles/virginia-leads-the-way-on-energy-affordability-in-2026/`. A companion measure, H.B. 1393, directs the incumbent provider to petition the Commission for rates recovering the cost of increased generating capacity from customers with demand of 25 megawatts or more.

4. What Virginia Did Not Do

It did not build a ceiling.

The tax prices consumption and does not limit it. A facility that doubles its draw pays twice as much and is not thereby out of compliance with anything, because there is nothing to comply with. The rate class assigns cost and sets no maximum. Nothing in the Commonwealth's 2026 enactments tells any data center that any quantity of electricity is too much, and the tax itself expires in two years.

So the score, from one state, in one session, is two of the three problems this Article identifies. Virginia built the instrument that Part III says is missing and gave the allocation doctrine of Part IV an operative number. It left the ceiling of Part II exactly where it found it. That is not a criticism of Virginia, which was writing a budget rather than an energy conservation program. It is the point. The two problems a state solved incidentally are the two that EPCA was never asked to solve, and the one it left standing is the one EPCA was built for and cannot reach.

V. What This Argument Does Not Show

An argument about invisibility is exposed to a particular failure: it can be answered by pointing at anything visible. This Part states the concessions in full, because each of them is true, each is available to a respondent, and none of them reaches the claim.

A. The Racks Are Full

Deployed accelerators are running at high utilization, and procurement lead times are long enough that relief from new supply is not near-term. Any suggestion that the industry has built ahead of demand in the way the fiber carriers did is not supported. The Part IV analogy is offered for the separation of the bet from the ticket, and that separation does not require overbuilding. It requires only that one layer be financed by shareholders and the other by ratepayers, which is true whether or not the demand appears.

B. Training Burns Power With Nobody Watching, and Nobody Counts It Here

A model in training consumes electricity at scale with zero users attached, by design. That consumption is outside every per-prompt figure discussed in Part III, and Google's paper says so in terms.[119] The figures also exclude image and video generation, which are substantially more energy-intensive per unit of output.[120] This cuts in favor of the Article's claim rather than against it: the published ratio is not merely missing its denominator, it is missing whole categories of numerator. It is also a caution against any reader who wants to reconstruct a total from these figures. It cannot be done from the outside, which is the claim.

C. Most Data Center Load Is Not AI

A majority of data center electricity consumption is conventional cloud computing: storage, web services, enterprise applications, the ordinary machinery of the internet. Attributing the whole of data center growth to artificial intelligence overstates the case and invites a correction that is easy to make. The Article's argument does not require the attribution. EPCA's blindness to servers is blindness to all of them, and it long predates the current models. Artificial intelligence is the occasion for noticing the defect and the reason its consequences have become large. It is not the defect.

D. Households Do Not Pay the Data Center's Electrons

A data center pays its tariffed rate for the energy it consumes. A residential customer's bill contains no line item for someone else's server. What a residential customer pays for is shared infrastructure, the capacity clearing price, and the delivered cost of the commodity, and the relationship between a large new load and any of those three is mediated by rate design, market structure, and resource adequacy in ways that vary by jurisdiction and cannot be assumed.

---

119. GOOGLE AI SCALE, *supra* note 5, at 5.

120. *See* AI INFERENCE BLOG, *supra* note 5 (limiting the estimates to text prompts); Butler, *supra* note 81.

This concession is the reason Part I.C.2 exists and the reason the empirical causation claim is disclaimed. An allocation argument that leans on an unproven causal attribution without saying so is a worse argument than the one this Article makes. The claim here is about who bears the cost of infrastructure capitalized to serve a load, under a doctrine that governs exactly that question, on a record that no outside party can construct.

E. The Doctrine Tolerates Crude Matching, and Presumes Network Benefit

The cost causation principle is more permissive than Part IV's framing alone would suggest, and two features of it cut against the argument. The Commission need not calculate benefits "to the last penny," and where it cannot quantify them but has "an articulable and plausible reason to believe that the benefits are at least roughly commensurate" with a class's share of regional sales, it may allocate on that basis.[121] The Commission may also "presume that new transmission lines benefit the entire network by reducing the likelihood or severity of outages."[122] A reader who expects Part IV to establish that generation built for data center load is misallocated will not find that holding here, and would not find it in the cases.

The concession is real and it narrows the claim to what Part IV.C asserts and no further. The doctrine's tolerance for crude matching presupposes that the crude inputs are available to both sides. The presumption of network benefit is rebuttable, and rebutting it requires the evidence of imbalance the doctrine demands from the challenger. Neither feature supplies the magnitude of the load, and neither is a reason the magnitude should be unavailable.

F. The Department Is Not the Villain

Nothing in Part II is a claim of agency bad faith. The Department applied the statute it has. It computed the household fraction because § 6292(b)(1)(B) told it to compute the household fraction, and the answer it got was the answer the arithmetic produces. Its 2014 retreat was

---

121. *Ill. Commerce Comm'n*, 576 F.3d at 477; *see also Ill. Commerce Comm'n v. FERC*, 721 F.3d 764 (7th Cir. 2013) (applying the standard and affirming).

122. *Ill. Commerce Comm'n*, 576 F.3d at 477.

a response to a jurisdictional objection that was correct on its own terms about the consumer program. Its construction of the computer room air conditioner category was a competent exercise of interpretive authority.

The point of setting the two episodes beside each other is not that the Department was inconsistent. It is that the two episodes together map the shape of the instrument: wherever the statute supplies a household denominator, the agency's reach ends at the front door of the house, and wherever it does not, the agency's reach extends to a category Congress never named.

VI. Objections

Part V stated what this Article does not show. This Part states who disagrees with what it does show, and why the disagreement does not dispose of it. The objections are drawn from a literature that is empirical rather than legal, and each is stated in its proponents' own terms before it is answered.

A. The Decoupling Objection

The strongest objection is that the premise is false. On this account the ratio and the total have not come apart at all, because efficiency gains have absorbed demand growth, and the assumption that rising computing demand must drive rising electricity use is precisely the error the literature has spent a decade correcting.

The objection has a distinguished pedigree. Shehabi, Smith, and Masanet modeled two decades of United States data center electricity demand and found that a rapid increase at the turn of the century had "significantly subsided to a nearly steady annual electricity use of about 70 billion kWh in recent years," and that "[w]hile data center workloads continue to grow exponentially, comparable increases in electricity demand have been avoided" through efficiency measures and consolidation into hyperscale providers.[123] Masanet and colleagues then took the argument to a general audience in SCIENCE, under a subtitle that states the thesis without

123. Shehabi, *Decoupling*, *supra* note 9, at 124030.

qualification: "Growth in energy use has slowed owing to efficiency gains that smart policies can help maintain in the near term."[124] Koomey and Masanet supplied the mechanism: "As long as there are still inefficient facilities to displace, large increases in demands for computing need not drive large increases in total electricity use."[125]

Three responses, in ascending order of importance.

First, the proponents hedged, and the hedge has come due. The decoupling paper's own abstract states that "the growth of data center electricity use beyond 2020 is uncertain, as modeled trends indicate that the efficiency measures of the past may not be enough for the data center workloads of the future," and that continued stabilization "will require new innovations in data center efficiency to further decouple electricity demand from the ever-growing demand for data center services."[126] The SCIENCE subtitle likewise limits the claim to "the near term."[127] These are forecasts with expiration dates attached by their authors, and the modeled period ended before the generative models arrived.

Second, the same authors have since revised. Masanet, Lei, and Koomey published in 2024 under the title *To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution*.[128] The title concedes both halves of this Article's factual predicate: the use is growing, and the data required to see it does not exist. The second concession is Part III's subject, reached independently by the authors of the objection.

Third, and dispositively: the objection does not reach the claim. Suppose the decoupling thesis is entirely correct and efficiency has absorbed demand growth without remainder. EPCA is entitled to none of the credit. The decoupling those authors document was produced by hyperscale consolidation, improved power usage effectiveness, and server utilization, and it was produced with no federal energy conservation standard applicable to any machine in the

124. Masanet, *Recalibrating*, *supra* note 9, at 984. The article is among the most cited in the field.

125. Koomey & Masanet, *supra* note 100, at 1625.

126. Shehabi, *Decoupling*, *supra* note 9, at 124030.

127. Masanet, *Recalibrating*, *supra* note 9, at 984.

128. Eric Masanet, Nuoa Lei & Jonathan Koomey, *To Better Understand AI's Growing Energy Use, Analysts Need a Data Revolution*, 8 JOULE 2427 (2024), `https://doi.org/10.1016/j.joule.2024.07.018`.

building, because no coverage determination for computer servers exists.[129] The Department withdrew the only one it ever proposed.[130] The mechanism Koomey and Masanet identify is a stock of inefficient facilities awaiting displacement, and a stock is a finite thing. When it is exhausted the mechanism stops, and nothing in the statute is positioned to notice, because the statute's instrument for noticing is a fraction with households in the denominator.[131] The decoupling literature describes an outcome with no statutory author and no statutory guarantor. Read carefully, it is the best evidence for this Article's thesis rather than against it: the one period in which data center electricity was arguably held in check is a period in which the federal efficiency program was not present.

B. The Plateau Objection

A second objection holds that the trend is bending rather than merely absorbed, and that measured reporting from inside the firms shows it.

Patterson and colleagues at Google argued in 2022 that four best practices reduce machine learning training energy "by up to 100x" and emissions "up to 1000x," that published estimates had overstated the real footprint by factors of "100x–100,000x," and that "the real-world energy usage trend of ML is fairly boring."[132] They predicted that if the field adopted those practices, emissions from training would decline by 2030.[133]

The reassurance offered in support of the prediction is the objection's undoing, and it is worth stating precisely. The paper reports that overall machine learning energy use at Google held steady below 15 percent of the company's total energy use across the preceding three years.[134] That figure is a ratio. Its denominator is Google's total energy use, and the sentence is silent on what the denominator did. Since 2019 the power consumed by Google's products

---

129. *See supra* Part II.D.

130. 79 Fed. Reg. 11,350, 11,350 (proposed Feb. 28, 2014); *see supra* Part II.D.

131. 42 U.S.C. § 6292(b)(2)(A).

132. Patterson et al., *supra* note 83, at 18. The author list includes Jeff Dean and Urs Hölzle.

133. *Id.* The prediction is the article's title.

134. *Id.*

and services has risen roughly 250 percent, and the company reported a 37 percent increase in electricity demand in the most recent year, which it describes as the largest load growth in its history.[135] A share that holds constant while the base triples is not a plateau. It is a proportion, doing what proportions do when nobody reports the number underneath them.

This Article does not raise that observation to convict anyone of bad faith. Patterson and colleagues disclosed the figure as a share because a share was what they had authority to publish, and they urged the field to publish more.[136] The point is structural and it is the point of Part III. The most sophisticated environmental accounting produced by the firm best equipped to produce it reported a quotient and left the divisor out, and the reassurance the quotient was offered to support did not survive contact with the divisor. Four years later the same company reported the largest load growth in its history.[137]

C. The Overestimation Objection

A third objection is methodological. Koomey and Masanet catalogued the recurring errors in assessments of information technology energy use, among them the practice of "making long-term projections" in a sector where "even projections extending only a few years are highly uncertain," and the application of exponential growth rates that "can result in eye-popping projected changes" and attendant media attention.[138] Their leading historical illustration is the one this Article uses for a different purpose: the 1990s extrapolation from Internet traffic "doubling every 100 days," which "led to vast overinvestment in network capacity around 2000."[139]

The objection is correct and it does not apply. This Article makes no projection. It does not estimate how much electricity artificial intelligence will consume in 2030, or in any year. It

---

135. GOOGLE 2026 REPORT BLOG, *supra* note 6; Clancy, *supra* note 7.

136. Patterson et al., *supra* note 83 (recommending that machine learning papers report emissions explicitly).

137. GOOGLE 2026 REPORT BLOG, *supra* note 6.

138. Koomey & Masanet, *supra* note 100, at 1625-26.

139. *Id.* at 1625; *compare supra* Part IV.A (drawing on the same episode for its financing structure rather than its forecasting error).

does not extrapolate a growth rate, and it takes no position on whether any published estimate is too high. Its claims are that the statute's coverage gate is a household fraction, that a firm published a quotient and withheld the divisor, and that a rate allocation doctrine assigns a burden to the party without the number. Each is a claim about a text or a document, and none is falsified by a forecast turning out wrong in either direction. Indeed the Koomey and Masanet critique and this Article point the same way. Their prescription is that analysts should stop extrapolating and start measuring. Part III observes that the measurement exists, inside one firm, and is not released.

D. The Rebound Objection

A fourth objection is that this is warmed-over Jevons, that the rebound literature has been arguing about these magnitudes for forty-five years without resolution, and that a legal article adds nothing by joining an unresolved economic dispute.

The premise is right and the conclusion does not follow. The magnitudes are contested and this Article contests nothing.[140] The reason is stated in the Introduction and repeated here because it is the whole distinction: set every rebound elasticity in the literature to zero and Part II is unchanged.[141] Rebound is a claim about how consumption responds to price. This Article's claim is that a jurisdictional gate cannot be satisfied by a machine that no household owns, which is a claim about statutory text and would be true in a world with no behavioral response at all. A reader who files this Article under rebound has filed it under the discipline that cannot fix it.

VII. Conclusion

The Energy Policy and Conservation Act is fifty-one years old and it has done what it was built to do. Refrigerators use a fraction of what they did. The program is, by the metric it was given, a durable success.

---

140. *See supra* note 17 (collecting the dispute).

141. *See supra* Part I.B.1.

The metric it was given is a ratio. Congress in 1975 wrote a statute that makes machines better and left out any provision that makes the country use less, and for most of the intervening period nobody had to notice, because the two moved together. They no longer move together. In the twelve months ending May 2025, the energy required to answer one median Gemini prompt fell by a factor of 33.[142] In the year that followed, the electricity demand of the company that achieved it rose 37 percent, the largest annual increase in its history.[143] Both numbers are true. The statute can read one of them.

The mechanism is a fraction. To bring a new consumer product under federal energy standards, the Secretary must find that its "average annual per-household energy use" is likely to exceed 100 kilowatt-hours, a quantity the statute defines as aggregate household energy use "divided by the number of such households which use products of such type."[144] The machines running the country's fastest-growing electrical load are in no household. They contribute to neither term. When the Department went looking for them in 2013, the statute made it ask how much electricity a server draws in somebody's house, and the Department answered 1,900 kilowatt-hours, and cleared the bar, and the racks in the data centers were not in the calculation at all.[145] Industry objected that the commercial units lay outside the consumer program, the Department agreed and withdrew,[146] and the commercial program's door, which has no numeric lock and opens onto the purpose of conserving "the energy resources of the Nation,"[147] has never been tried for a server. It was tried for the air conditioner in the same room, which the Department covered by holding that a 1978 definition reached a category "which EPCA did not expressly identify."[148]

---

142. GOOGLE AI SCALE, *supra* note 5, at 1.
143. GOOGLE 2026 REPORT BLOG, *supra* note 6.
144. 42 U.S.C. § 6292(b)(1)(B), (b)(2)(A).
145. Computer Systems Proposal, *supra* note 51, at 11,347.
146. 79 Fed. Reg. at 11,350.
147. 42 U.S.C. § 6312(a)-(b).
148. CRAC Standards Rule, *supra* note 60, at 36,396.

The disclosure regime has independently built the same instrument. Google measured the number of prompts, divided by it, published the quotient, and withheld the divisor, and when asked for it said the company was "not comfortable revealing that for various reasons" and wanted users thinking about the energy per prompt.[149] OpenAI has published no model-specific figure in six years,[150] and outside researchers estimating GPT-5 had to reconstruct the numerator from hardware assumptions and take the denominator from press reports.[151] A statute written for an oil embargo and a technical report written by an engineering organization fifty years later arrived at the same shape, from opposite directions, because a ratio is the only environmental number an institution can improve forever.

And so the allocation question arrives with its evidence held by the parties it would be asked about. Cost causation asks whether the customers billed for a facility derive benefits roughly commensurate with the cost.[152] Answering it requires the magnitude of the load. The magnitude of the load is the product of a published ratio and a confidential count.

There is no ceiling, because Congress never wrote one and no agency can supply it. There is no denominator, because the only parties who have it have said they would rather not. And there is no exit, because the layer that carries the cost is the one layer of this economy that a household cannot leave.

149. Crownhart, *supra* note 13.

150. COMPUTING, *supra* note 90.

151. Morales, *supra* note 92; DIGITIMES, *supra* note 93.

152. *Ill. Commerce Comm'n*, 576 F.3d at 476.