

Allocating the AI Kilowatt-Hour: Why the Energy Market Is Not a Protocol

From Regulating Models to Making Material Arbitrations Explicit.
Second Volume of a Diptych on AI's Physical Constraints

I. A Threshold Crossed: the Constrained Resource Is Now Being Allocated

The article that opens this diptych, [*The Supposed Virtuality of AI Against the Wall of the Real \(February 2026\)*](#), established that generative AI has entered a heavy-industry regime, constrained by a converging set of finite resources: energy, water, silicon, data. Three months later, the diagnosis is being confirmed empirically faster than expected.

The International Energy Agency now projects 1,100 TWh of datacenter consumption by the end of 2026 (an 18% upward revision from the December 2025 estimate), with roughly 5 to 15% attributable to AI today and 35 to 50% anticipated by 2030. In the United States, *Fortune* (April 20, 2026) documents that these datacenters absorb on their own about half of the *growth* in national electricity consumption; retail prices have risen by some 40% over 2021-2026; Data Center Watch records, for the period from May 2024 to March 2025, sixteen projects affected by local opposition (six fully blocked, ten delayed) for \$64 billion, a figure to which Q2 2025 alone adds twenty more projects for \$98 billion, with the 2025 cumulative reaching at least forty-eight projects for a total volume well above \$100 billion. Goldman Sachs estimates a 0.1 percentage-point impact on US core inflation in 2026 and 2027.

These figures are not forecasts. They are arbitrations already in motion.

The relevant question is therefore no longer: *is AI consuming too much?* It is: *by what mechanism is this consumption being distributed across competing uses?* The object has shifted. Materiality is no longer the hypothesis to be demonstrated; it has become the premise of the arbitration. And the current arbitration is triply distributed (energy market, long-term contracts, local opposition), without any of these three mechanisms being designed to deliberate.

Scope of validity: the analysis that follows draws principally on data and cases from the United States and the European Union. Transposition to other regions (East Asia, Latin America, Africa) calls for a contextual update of its own, which this article does not claim to provide.

II. Where the Decision to Allocate an MWh to AI Is Made Today

Three Levels to Separate

Before any cartography, an analytical distinction is required. Three levels are frequently conflated in public debate, and the conflation blocks the argument:

- *Physical allocation*: flows of energy actually delivered to a use (kWh transiting from the grid to the datacenter).
- *Economic allocation*: the price signal that ranks capacities to pay for that resource (€/MWh, contracts, premiums).
- *Normative arbitration*: the collective decision regarding the relative social value of the uses to which the resource is devoted.

Physical allocation happens in real time and poses no problem: the grid ensures the physical delivery of the requested power. Economic allocation operates through the energy market and works on its own dimension. Normative arbitration, on the other hand, is executed by no dispositive at all. The implicit slide by which the three levels are equated (*the market allocates, therefore the market arbitrates*) is precisely the dominant fallacy. Price aggregates solvent preferences under local constraint; it does not aggregate non-solvent collective preferences, deferred externalities, or differentiated social returns. This proposition shifts the discussion from normative judgment to mechanical description: it does not judge markets; it describes what they do and do not do.

Cartography in Three Layers

The energy market. First layer. The price signal locally aggregates information on scarcity, forms a price that circulates among producers, transmission operators and consumers, and allows continuous reallocation. This mechanism functions on its own dimension, economic allocation. It nonetheless presupposes two conditions in order to incidentally produce an acceptable normative arbitration: that the actors competing for the resource pay an effectively comparable marginal price, and that their capacity to pay correlates with the social value of their uses. For AI datacenters, the first condition is only nominally met; the second cannot be presumed.

Long-term contracts. Second layer. Power Purchase Agreements (PPAs) signed over horizons of fifteen to twenty years remove durable volumes from the spot mechanism. Microsoft has contracted the recommissioning of Three Mile Island Unit 1 with Constellation Energy (around 835 MW), plus more than 10 GW of additional renewable capacity by 2030. Amazon Web Services aims for over 5 GW powered by SMRs by 2039. Google has signed a 500 MW contract with Kairos Power and publicly commits to a 24/7

carbon-free supply. These contracts are not a market distortion; they are the rational use of the market by actors who can afford the security premium. Their net effect, however, is one of *subtraction*: the volumes thus engaged are withdrawn from the spot mechanism for the duration of the contract. Contractual time horizons exceed electoral time horizons by a factor of three to four.

To this private subtraction, the regulatory frameworks in force oppose no procedure for use-arbitration. The European AI Act, whose enforcement powers over GPAI models will activate on August 2, 2026, contains no provision for the hierarchization of uses on a resource criterion. The NIST AI RMF Profile *Critical Infrastructure*, published as a concept note on April 7, 2026, addresses risk management for critical-infrastructure operators, not the public stratification of their AI uses.

Local blockages. Third layer. Forty-eight projects blocked or delayed in 2025 in the United States according to Data Center Watch, for a volume well above \$100 billion; the Irish cap on the datacenter share of national consumption; moratoria in Singapore and the Netherlands on new projects; controversies over agricultural mega-basins in France, already mobilized as a procedural analogy in the parent article. These oppositions produce a real but unpredictable regulatory effect. Their logic is not arbitration between uses: they do not say *to what* the contested project should have yielded, but oppose a veto on a particular use in a particular location. It is an *ex post sanction*, not an *ex ante arbitration*.

Mechanistic Demonstration of the Blindness

The energy market's blindness to social returns is not an accidental defect. It is mechanical, and it is so in this case for three reasons that compose a closed causal system.

First, *non-differentiation by end use*. Electricity is fungible at the point of delivery. The grid does not know, by construction, how to distinguish a MWh devoted to an MRI scanner from a MWh devoted to a recreational conversational agent: it delivers electrons to the meter. This fungibility is technically necessary; its direct consequence is that any differentiation by end use falls to the institutional dispositive, not the technical one.

Second, *capital asymmetry*. The capacity to secure long-term energy access depends on the capital volume of the actors in competition. A hyperscaler commits billions in energy CapEx; a midsize hospital, an academic laboratory, a critical-infrastructure operator that is not solvent at the same standard cannot deploy these arrangements.

Third, *time horizon*. Contracts of fifteen to twenty years exceed the time horizon of the spot signal. Once engaged, they durably remove the most strategic volumes from the market. The price signal continues to function, but on a subset of the resource, the part that has not been pre-allocated.

The formal proposition that follows fits in one sentence: *in a system where the resource is fungible at the point of delivery, capital is concentrated, and intertemporal commitments are dominant, the price signal converges structurally toward allocation by solvency and not by social value.* This is not a critique of markets in general. It is a description of what this market does, here, under these specific conditions.

We have built a metric for functional aptitude. We have never built a metric for socio-technical efficiency, and that is why allocation follows solvency.

Three Emergent Properties of the System

The market-contracts-blockages system is more than a cartography. It exhibits three emergent negative properties that are observable and nameable: *absence of collective reversibility*, the volumes pre-allocated by long-term contracts exceed the time horizon of any ordinary political dispositive; *absence of comparability of uses*, the price mechanism equalizes solvencies, not social values; *absence of an institutional point of entry*, no mechanism explicitly designed exists to arbitrate, upstream, the allocation of the resource between competing uses.

The formula that summarizes this distributed system fits in a single sentence: **the market allocates, contracts lock in, oppositions sanction; none of the three deliberates.** Allocation exists. Procedure does not. The system is stable, unmandated, not collectively revisable, and not informed by any public metric.

III. Failure of the Instrument: the Footprint-Free Benchmark

The procedural failure is bound to a metrological failure. The bridge between the two is not analogical. It is necessary.

A decision system can only arbitrate what it measures. In the absence of a metric coupling performance and resource cost, any allocation is mechanically decoupled from value. The proposition does not depend on a school of economics or an ideological option. It describes a condition for instrumenting judgment. For an arbitration between uses to be possible, one must have a metric that couples what a use produces to what it consumes. This metric does not exist, in the sense that it has never been built as a public infrastructure, and in the sense that the benchmarks acting as de facto standards are precisely defined to measure performance *independently* of resource cost.

This is what the Twingital corpus designates, under the name RAISE, as the critique of the *footprint-free benchmark*. An MMLU at 89%, an SWE-bench Verified at 84%, a GAIA at 65% describe functional aptitude; they say nothing about the energy, water, amortized silicon, or GPU time mobilized to reach that aptitude. This silence is not neutral. It makes

performances comparable solely on functional accuracy, and structurally non-comparable on the dimension that matters for arbitration: the resource yield for a given task.

On April 12, 2026, the Center for Responsible Decentralized Intelligence at Berkeley published an empirical demonstration that the eight principal agent benchmarks (SWE-bench, WebArena, OSWorld, GAIA, Terminal-Bench, FieldWorkArena, CAR-bench and one further) could be systematically saturated by automated reward hacking, achieving near-perfect scores without solving a single real task. The study covered eight existing benchmarks; it demonstrates the logical fragility of a *class*, the class that does not incorporate resource cost, rather than a universal empirical vulnerability, but the underlying mechanism applies to any metric designed in the same way. A benchmark that does not constrain its optimization with a resource cost is, by construction, exposed to all the strategies that maximize the score without producing the corresponding resolution effort under operational conditions. The fragility observed is the foreseeable consequence of a metric designed to be maximized rather than to be informative.

This fragility of agent evaluation does not, on its own, prove an inability to arbitrate an MWh. It becomes such an inability when these benchmarks step outside their evaluation function and enter an allocation function, that is, when they become investment signals, public-procurement criteria, or sectoral certification bases. The bridge between metrological failure and allocative failure is not conceptual. It is institutional. The slide by which a technical test transforms into an admission key, without having been designed to bear that load, is what the Twingital corpus designates under the name *promotion gateway* (in the original French: *port de promotion*) of the benchmark. The metric makes a career beyond its jurisdiction of origin, and its use within a decisional framework for which it was not calibrated produces foreseeable distortions.

The Stanford HAI AI Index 2026 documents another symptom. The hallucination rates of the 26 frontier LLMs evaluated range from 22% to 94% depending on the domain; generalist benchmarks of the MMLU and MMLU-Pro type are functionally saturated above 88%, which makes between-model gaps at those scores statistically unexploitable. Gartner indicates, for its part, that only 28% of AI projects in infrastructure and operations deliver the promised ROI, with one in five projects collapsing during execution. The capacity metric correlates weakly, or does not observably correlate, with the metric of operational value; and neither of the two is instrumented with a resource-cost metric.

The precedent lies elsewhere. The financial crises of the 2000s did not stem from an absence of market. They stemmed from an institutionalized and inadequate risk metric (VaR, AAA ratings on subprimes) whose failure contaminated allocation across the system. In healthcare, decades of overtreatment did not arise from a defect of clinical evaluation. They arose from a selection of primary endpoints (intermediate criteria insufficiently linked to final clinical benefits) whose inadequacy oriented downstream

medical decisions. The problem, in both cases, is not the absence of the instrument. It is its *inadequacy*. And it is this same inadequacy that operates in the current AI ecosystem.

Metrological blindness and political blindness share the same structural cause. The construction of the instrument and the construction of the decisional dispositive have in common that both were carried out outside any requirement of resource cost. To correct one without the other is to preserve the failure.

IV. Three Structural Failures of the Implicit System

The cartography of section II describes a system that *works*. It does effectively allocate scarce resources to uses, without fatal interruption, in the presence of major shocks: a turn in US public opinion, local blockages exceeding \$100 billion. This functional robustness is precisely what makes the critique difficult: energy markets are not broken, long-term contracts are not unlawful, local oppositions are not pathological. The system does not malfunction on its own dimension, technical and economic. It decides on allocation, but cannot arbitrate the normative question.

Three failures shed light on the single underlying cause.

Non-comparability of Uses

A megawatt-hour is not homogeneous to a megawatt-hour. A MWh devoted to modeling North Atlantic climate trajectories, a MWh devoted to a diagnostic-assistance device in pediatric oncology, and a MWh devoted to a recreational conversational agent do not carry the same social returns. The energy market does not draw this distinction. It cannot, by construction, the three mechanisms of section II locking it down. The market does not compare uses, it compares solvencies. When the capacity to pay correlates with social value, the mechanism incidentally produces an acceptable arbitration. When the two diverge (and the divergence is precisely the contemporary situation, where heavily capitalized actors can pay for uses whose marginal social value is low), the market allocates in a manner *Pareto-efficient yet mute on value*.

Intertemporal Lock-In

Long-term contracts, economically rational from the standpoint of the hyperscaler whose margin depends on guaranteed access to several pilotable gigawatts, produce a lasting subtraction from the spot market. The issue is neither the legality nor the efficiency of the mechanism: it is the asymmetry of time horizons. Ten years of legislative mandate, five years of executive mandate, two to three years of regulatory cycle, set against fifteen to twenty years of PPA. The geometry is asymmetric, and that asymmetry is itself a political fact. The window for revision before contractual lock-in closes between 2027 and 2028. Beyond that threshold, fifteen to twenty years of hyperscaler PPAs dominating the marginal volumes would reduce the public arbitration capacity to almost nothing, not

because the contracts would forbid it, but because they would have preempted the matter of the arbitration.

Metric Substitution: the "Green PPA"

The parent article established at paragraph 8.4 that a renewable PPA can, when physical additionality and temporal concordance are not demonstrated, operate essentially as a *bookkeeping reallocation* of existing production. Not all PPAs fall into this case: some do effectively finance bas-carbon capacity that would not have been built otherwise. But when the conditions below are not met, the risk is to displace legitimacy toward an attribution-accounting, regardless of the physical reality of the electron flow. Three precise technical properties anchor the critique.

Physical non-additionality: when a PPA is not associated with the construction of bas-carbon capacity that would not have existed without it, the purchase of guarantees of origin attributes to the buyer, in accounting terms, renewable production available elsewhere; the residual fossil production continues to cover net demand elsewhere in the system.

Temporal desynchronization production / consumption: datacenters consume 24 hours a day; intermittent renewable sources (solar, wind) produce according to variable, uncorrelated profiles. Without long-duration storage capacity at scale, annual or even monthly matching masks an hourly imbalance that pilotable fossil production continues to fill.

Illusion of energy traceability: the bookkeeping attribution of a certificate to a consumption creates a representation of a clean energy chain (*24/7 carbon-free, 100% renewable matched*) to which no direct and continuous physical relation, in the physical system, between attributed production and effective consumption corresponds. The metric is correct on its own bookkeeping perimeter; it is misleading on the physical dimension it claims to summarize.

These three properties are admitted in the technical literature. Their effect, recast within the present article, is cumulative: the booked decarbonization metric *takes the place of debate* on allocation. It shifts the public conversation from a terrain where material cost would be explicit to a terrain where environmental performance, even partially fictional, occupies the space of legitimacy. We have invented *24/7 carbon-free* before inventing *24/7 socially-justified*.

Three failures, one underlying cause: the absence of a framework linking resource, use, and value. The market-contract-blockage system is stable and uninstructed. This combination is not indefinitely sustainable. US public opinion turns as retail prices rise and blockages multiply; territorial jurisdictions improvise caps whose normative value is weak but whose signaling value is strong; long-term contractual lock-in expands as actors

anticipate possible regulation. These are not three independent movements. They are three coherent consequences of the same cause.

V. Conditions for a Protocol: Differentiated Conditions of Access, Not Planning

The libertarian objection must be addressed head-on. To build a sectoral deliberative protocol around the allocation of the AI kilowatt-hour is to open the door to planning nostalgia, political capture, arbitrary sluggishness. It is, more brutally put, to trade a distributed and efficient decision for a centralized and contestable one.

The objection is serious. It is also, in its current form, miscalibrated. It assumes that the alternative to a protocol is an ideal competitive market. But the alternative to a protocol is the captive market described in section II: a price signal bounded by long-term contracts and corrected by NIMBY opposition. This is not the canonical configuration that theories of market efficiency mobilize to defend the supremacy of the price signal. To reframe the debate in terms of protocol therefore amounts to comparing not *market vs. planning*, but *captive market without procedure vs. market constrained by procedure*. The protocol does not substitute for the price mechanism. It frames it.

Doctrinal lock-in of the distinction: **a protocol for AI allocation does not classify uses to assign them the resource directly. It defines differentiated conditions of access to the resource.** This precision shifts the object of the dispositive from a logic of central allocation (which no reasonable proposal advances) to a logic of controlled admissibility (which existing sectoral arrangements have practiced for decades). Concretely, differentiated conditions of access can take the form of stronger evidentiary requirements for certain uses, of minimum efficiency constraints variable by stratum, of asymmetric transparency obligations according to externality profile. None of these forms amounts to prescribing who consumes what.

Four partial precedents, mobilized as reference points and not as transposable models, illuminate the conditions of practicability. The allocation of radio spectrum bands by the ITU and, in France, by ARCEP, organizes differentiated conditions of access to a scarce resource with periodic revisions, without paralyzing industrial innovation. The French SDAGEs (Schémas Directeurs d'Aménagement et de Gestion des Eaux, Master Plans for Water Development and Management) articulate European, national and basin levels to hierarchize water use under constraint. The management of airport slots combines price signal and admissibility constraints in a sector where the scarcity of the resource is non-substitutable. The French Citizens' Climate Convention of 2019-2020 is to be mobilized not as a naive inspiration but as a procedural cautionary tale: its limited application rate shows that deliberative procedure is legitimate *only* when it is articulated to a clear executive mandate and to observable monitoring mechanisms.

None of these arrangements is transposable as is. All indicate that the idea of a sectoral protocol is in no way utopian: it has, in fact, more functional historical precedents than the current alternative.

Four conditions appear necessary to the practicability of such a protocol.

1. **Condition 1, performance / footprint coupling.** Any performance reported by a benchmark used as an institutional gate (compliance audit, public procurement, sectoral certification) must be accompanied by a normalized resource cost. This requirement is the direct consequence of the RAISE critique of the footprint-free benchmark. **Without a coupled metric, the protocol is impossible. With a coupled metric, the market becomes arbitrable.** The protocol does not replace the market; it makes it operational on a dimension it does not currently see. This inversion is the conceptual pivot of the entire proposal. Once the metric is available as public infrastructure, arbitrations can again partly proceed via the price signal, on dimensions that aggregate performance and resource cost, which is precisely what none of the three current mechanisms aggregates.
2. **Condition 2, explicit differentiated conditions of access.** Distinguish publicly, through a legible dispositive, strata of use (health and critical infrastructure, scientific research and modeling with strong positive externality, productive automation, agentic and non-productive inference) to which distinct conditions of access apply: requirements of social-utility evidence, minimum efficiency constraints, operational transparency obligations. This stratification is not planning. It is the explicitation of arbitrations already implicitly made by hyperscalers, local jurisdictions and grid operators. The shift from implicit to explicit is the political gesture proper.
3. **Condition 3, multi-jurisdictional level.** European scale for the grid of uses, prevention of carbon leakage, and the minimal harmonization that prevents a race to the bottom; national scale for territorial implementation and infrastructure planning; local scale for acceptability and offsets. Without this articulation, one falls either into a European harmonization disconnected from the field, or into contradictory local decisions. A protocol unilaterally adopted by a single jurisdiction (European, for instance) would inevitably produce a *carbon / regulatory leakage*: displacement of the most resource-intensive AI workloads toward jurisdictions without a protocol, amplified geographic concentration, deepening of asymmetries between blocs. Multi-jurisdictionality is not a refinement; it is a condition of non-circumvention.
4. **Condition 4, institutional reversibility.** Five-year revision cycles, observation of unanticipated effects, possibility of redefining the strata and the conditions of access. Without programmed reversibility, the protocol reproduces the lock-in it claims to correct. The contractual time horizon of hyperscaler PPAs (15 to 20

years) must be absorbed by a shorter revision time horizon, failing which the procedure itself becomes an instrument of subtraction.

At an operational scale, operational projects of AI allocation in critical contexts (pediatric healthcare in particular) demonstrate that a stratification of uses is technically operable when high-criticality inference loops, predictive workloads and documentary tasks are not placed in the same admissibility class, each class carrying its own requirements of evidence, transparency and efficiency. This internal operability, observed in regulated contexts (pediatric medicine, dispositives under HDS constraint and ISO compliance), does not constitute a general proof of the proposed sectoral protocol. It does, however, refute, by existence, the objection that a stratification of uses would be a conceptual object beyond operational reach. It *is* possible, within a constrained and regulated system, to define differentiated conditions of access that do not paralyze production. The passage from internal operability to a public sectoral dispositive is not trivial: it is the proper institutional work that European regulators will have to conduct in the 2026-2028 window, before the lock-in by long-term contracts closes the arbitration window.

VI. Conclusion: from Regulation to Allocation

The parent article concluded on a question: what share of our finite physical resources will our societies accept to allocate to AI, and at the expense of which other uses. The present article does not answer that question. It shifts its object.

AI is no longer merely an object to regulate. It is a vehicle of allocation of scarce physical resources. To regulate its models, its high-risk uses, its biases: the AI Act does it, the British AISI evaluates them, NIST builds the standards. None of these instruments decides, and none has the vocation to decide, how the kilowatt-hours, the hectoliters and the hectares devoted to this industry are arbitrated against their alternative uses. The regulation of models does not stand in for the arbitration of material allocation. The diptych formed by these two articles supports a simple thesis: doctrine must explicitly separate the two, and instrument each through its proper dispositives.

The real cost of AI is not measured solely by the energy it consumes. It is also measured by the trajectories that this energy makes impossible before they have been deliberated. A society can always burn a scarce resource to produce an immediate benefit. The political question begins when that same resource could have opened a longer trajectory, and the decision to forgo that trajectory has never been formulated as a decision.

What the current situation conceals, as it stabilizes, is a major political decision taken without political procedure. Microsoft, AWS, Google and their peers decide, through their long-term PPAs, the energy trajectory of the next two decades. Local jurisdictions correct, by way of moratoria and blockages, what they never debated upstream. US public opinion turns without having been consulted. The system functions and erodes simultaneously.

A society does not choose whether it pays the material cost of AI. It chooses whether that cost is borne by default or distributed according to explicit criteria. **We have industrialized the production of artificial intelligence. We have not yet industrialized the decision of which uses deserve to consume its matter.**

References

- Parent article: Vetillard, J. (2026). *La virtualité supposée de l'IA face au mur du réel* [The Supposed Virtuality of AI Against the Wall of the Real]. Twingital Institute, revised edition February 2026. <https://twingital-ventures.com/fr/publications/ia-energie-contraintes-physiques/>
- International Energy Agency (2026). *Energy and AI / Global Energy Review 2026*. Projection 1,100 TWh datacenters in 2026, AI share 5-15% today / 35-50% by 2030.
- *Fortune* (April 20, 2026). "Data centers now account for half of all new U.S. electricity use, just as Americans start to sour on AI."
- Data Center Watch (initial report May 2024-March 2025; Q2 2025 update). 16 projects / \$64 billion for the first period; +20 projects / \$98 billion in Q2 2025 alone; 48+ projects / >\$100 billion for the 2025 cumulative.
- Goldman Sachs Research (February 2026). Note on the impact of AI datacenters on inflation. +0.1 pp core inflation US 2026 and 2027.
- UC Berkeley Center for Responsible Decentralized Intelligence (April 12, 2026). *Trustworthy benchmarks (cont.)*.
- Stanford HAI (2026). *AI Index Report 2026*.
- Gartner (April 7, 2026). "AI projects in I&O stall ahead of meaningful ROI returns."
- NIST (April 7, 2026). *Concept note for AI RMF Profile on Trustworthy AI in Critical Infrastructure*.
- European Parliament (2024). Regulation (EU) 2024/1689 (AI Act).
- Sorrell, S. (2009). "Jevons' Paradox revisited." *Energy Policy*, 37(4).
- Kaack, L. H. et al. (2022). "Aligning AI with climate change mitigation." *Nature Climate Change*, 12.