A Common Electrical-Energy Basis for Comparing Lunar Oxygen and Propellant Production Routes

Walter Kueffer

Draft manuscript, version 0.13 (2026-07-01)

Code, data, and all figures are reproducible from the open-source model at https://github.com/dubthree/lunar-propellant-energy-model.

Abstract

Lunar in-situ propellant production is widely proposed as the key to affordable cislunar transport, yet the basic engineering question (which oxygen-extraction route costs the least *energy* per kilogram of product) is hard to answer from the literature, because published figures are reported on incompatible bases: hydrogen reduction as electrical energy, carbothermal reduction as thermal energy, and molten regolith electrolysis (MRE) with no like-for-like electrical figure (NASA reactor-sizing models report whole-system numbers at much wider boundaries). Prior work compares subsets (Taylor and Carrier (1993) reviewed ∼20 processes; Leger et al. (2025) modeled several on a common basis), but no published comparison places all five of these routes on a single electrical-equivalent basis with paired uncertainty propagation; that common basis, not any first-ever number, is this paper's contribution. We present a small, open, uncertainty-quantified model that places five routes (hydrogen reduction, carbothermal reduction, MRE, molten-salt electrolysis, and polar water-ice mining) on one electrical-equivalent kWh/kg O_2 basis under a single explicit system boundary, with Monte-Carlo propagation of literature parameter ranges. Continuous standing losses are charged symmetrically: every high-temperature route carries a reactor heat-loss term grounded in the only measured carbothermal datum (NASA CaRD), and the low-temperature water route carries its own charge sheet (vapor-capture efficiency, cryogenic excavation, a permanently-shadowed-region standing loss, and the ice's full thermal chain), so no route is exempted from the loss categories that burden its competitors. We check the framework against the one route with a clean published electrical figure (hydrogen reduction, Leger et al. (2025), 24.3 ± 5.8 kWh/kg LOX): our loss-free configuration of that route reproduces Leger's central value (nominal 24.3, Monte-Carlo median ∼21), reported as substantial interval overlap between two independent estimates. Using a paired Monte Carlo that shares common parameters across routes, we report dominance probabilities rather than overlapping error bars. The central finding is that **on an all-electric basis the polar water-ice route is the most energy-efficient (cheapest in 89% of paired trials) and the only route that yields a complete LOX+LH_2 propellant**, though its margin narrows once it is charged for capture, excavation, and standing loss: water (∼16 kWh/kg O_2) versus a high-temperature cluster at ∼21–32, with hydrogen reduction the most likely worst once its standing loss is charged. The advantage is structural: water mining is the only low-temperature route (sublimation at ∼273 K), so it escapes the large, continuous reactor heat-loss penalty (∼63–93 kWh-thermal/kg O_2 in the one measured carbothermal demonstration) that burdens every high-temperature thermochemical or electrochemical route. The result is, however, conditional on heat being electrically supplied: a solar-thermal sensitivity shows that concentrated sunlight at a sunlit site inverts the ranking (carbothermal falls to ∼8 kWh-electric/kg O_2), at the cost of concentrator mass and of being at the wrong site for water. A sensitivity analysis identifies the single highest-value measurement for each route, we translate energy into required surface power and landed fission-plant mass, and we outline two companion analyses on reusing co-located compute waste heat (which is, fittingly, most useful precisely at the low-temperature water route).

1 Introduction

The case for lunar propellant rests on energy: every process that turns regolith or polar ice into liquid oxygen (and ideally liquid hydrogen) is energy-intensive, and surface power is the binding constraint on any near-term architecture. Yet the comparison that should drive route selection has never been made on equal footing. Three incompatibilities recur in the literature:

1. **Thermal vs electrical energy are conflated.** Carbothermal reduction is reported as delivered thermal energy or "g O_2/kWh thermal"; hydrogen reduction is reported as electrical kWh/kg. On the Moon both ultimately draw on the same scarce electrical supply, but they are not interchangeable at face value.
2. **System boundaries differ.** Some figures include excavation, beneficiation, and liquefaction; others count only the reactor.
3. **Two routes have no like-for-like figure.** Molten regolith electrolysis and molten-salt electrolysis are described qualitatively or at incompatible whole-system boundaries; no published number for either sits on the same basis as the reduction routes.

The consequence is that capital allocation, power-plant sizing, and architecture choice rest on a comparison that does not exist. This is a modeling gap, not a hardware gap, which is precisely why it can be closed now, cheaply, and reproducibly. We build that comparison and quantify its uncertainty.

2 Methods

2.1 Functional unit and system boundary

We compute the **electrical-equivalent energy in kWh per kg of O_2 delivered to cryogenic storage**, and additionally report kWh per kg of total propellant for the one route that co-produces hydrogen. The system boundary is a fixed sequence of stages; every route enables a subset, using the same stage sub-models, so all differences between routes arise from parameters rather than inconsistent accounting:

> excavation/acquisition → beneficiation → heating (sensible, plus fusion for melt routes) → reaction (reduction enthalpy or Faradaic electrolysis) → product cleanup → water electrolysis (where the route produces H_2O) → gas compression → liquefaction (LOX always; LH_2 where hydrogen is retained).

2.2 Thermal-to-electrical conversion

To compare a thermally-reported route against electrically-reported ones, all thermal demand is converted to electrical-equivalent through an explicit electric-to-thermal efficiency parameter (resistive heating, nominal 0.90). This single, visible knob is what makes the comparison honest. The choice of an all-electric baseline is itself a siting assumption: it is exact for a permanently shadowed region and for nuclear-powered architectures, but at a sunlit site concentrated sunlight can supply high-grade process heat directly. Because that choice turns out to matter to the ranking, we report a solar-thermal pathway as an explicit sensitivity (Section 5), not bury it.

2.3 Uncertainty propagation

Every uncertain input is a triangular `(low, nominal, high)` distribution sourced from the literature (the full parameter table, with a citation on every value, is in `params.py`). A 20,000-trial Monte Carlo propagates these to a 90% interval per route. We propagate stated literature ranges, not assumed Gaussian measurement error. The electrolysis cell voltage and current efficiency of the electrochemical routes are sampled with a physical anti-correlation (a shared operating-severity latent: higher current density raises voltage and lowers efficiency together), avoiding unphysical parameter corners. The water route's new charge-sheet parameters (capture efficiency, cryogenic excavation specific energy, PSR standing loss; Section 2.6) are sampled the same way.

2.4 Ranking by paired Monte Carlo

Reading a ranking from five marginal error bars is a statistical error, because the routes share parameters (regolith specific heat, heat recuperation, electrolysis efficiency, liquefaction) that must take the same value in any given world. We therefore run a paired Monte Carlo: one shared parameter draw per trial across all routes, evaluated trial by trial, so we can report P(route A cheaper than route B) and P(route is cheapest / worst).

2.5 Continuous standing losses, charged symmetrically

A reactor held continuously at 800–1900 C radiates and conducts heat to its surroundings independently of the per-kg sensible heating, a term the sensible-heat-only stages omit. The only measured carbothermal datum (NASA's CaRD brassboard) implies ∼63–93 kWh-thermal per kg O_2 for the reduction step, dominated by this loss at demonstrated (tiny) scale. A scaled, well-insulated plant amortizes it over far more throughput, so the true value is highly scale-dependent. We therefore add a single reactor-loss term (log-triangular, nominal 8, range 2–30 kWh/kg O_2: 2 for a large well-insulated plant, ∼30 approaching the brassboard upper bound) to **all four** high-temperature routes, hydrogen reduction included. An earlier version of this model exempted hydrogen reduction on the argument that its Leger 2025 anchor already reflected realistic reduction energy; that exemption was inconsistent with the model's own claim that the route is built bottom-up from independent parameters, and it undercharged the route by ∼8 kWh/kg at nominal. The headline now charges the loss symmetrically; the loss-free configuration is retained as an explicit model option and is what the Leger validation tests against (Section 3.1). This term, omitted entirely in an early version of the model, is what once made carbothermal appear cheapest.

The water route is not exempt from standing loss either. Its sublimation zone runs at ∼273 K inside a 40–110 K environment, so it radiates and conducts heat into cold regolith that yields no water. We charge it a PSR standing-loss term (log-triangular, nominal 1.5, range 0.3–8 kWh/kg O_2), smaller than the reactor term because radiative loss scales as T^4 and 273 K is far below 1100–1900 C, but not zero.

2.6 The water route's charge sheet

An earlier version of this model charged the water route for thermal mining, electrolysis, and liquefaction and nothing else, which exempted it from every loss category that burdened its competitors. Version 0.13 charges it symmetrically:

- **Vapor-capture efficiency** (nominal 0.75, range 0.50–0.95): tent or cold-trap capture over fissured terrain in vacuum does not collect every sublimated molecule; mining heat and regolith throughput scale by its inverse.
- **Cryogenic excavation** (nominal 30 kJ/kg regolith, log range 2–200): icy regolith at 40–110 K has concrete-like strength; rock-cutting specific energies are orders of magnitude above the dry-simulant bucket-drum figure used for the regolith routes. The wide log range spans radiant in-situ concepts (low end) to mechanical cutting of competent icy ground (high end).

- **PSR standing loss** (Section 2.5).
- **The ice's own thermal chain**: sensible heat of the ice from feed temperature to 273 K ($c_{p,\text{ice}}$ ~1.75 kJ/kgK) and a reconditioning term (nominal 0.12 kWh/kg O_2) for refreezing captured vapor at the cold trap and re-melting it for the electrolyzer feed.
- **One correction in the route's favor**: the regolith sensible term now uses a cryo-range specific heat (nominal 0.45 kJ/kgK over 70–273 K) instead of the 1.15 enthalpy-mean over 250–1300 K, which had overstated the water route's sensible heat roughly twofold. The hot routes keep the hot-range value.

The net effect of the charge sheet is to raise the water route's nominal from 14.0 to 15.8 kWh/kg O_2 and to widen its interval; the route survives it (Section 4), which is a far stronger statement than the earlier result that exempted it.

3 Validation

3.1 Hydrogen reduction against Leger 2025 (and a second anchor)

Hydrogen reduction is the only route with a clean published electrical figure: 24.3 ± 5.8 kWh/kg LOX, with the reduction step ~55% and water electrolysis ~38% of the total (Leger et al., 2025). We built our estimate bottom-up from first principles and independent parameters; the largest tunable term (water-electrolysis efficiency) is set from an independent SOEC source (Hauch et al., 2020; International Energy Agency, 2019), not from Leger's implied value.

The comparison is made in the model's **loss-free configuration** (continuous standing loss zeroed), because Leger's full-chain figure contains no continuous standing-loss term that we can identify; the agreement below is evidence that his boundary, like our loss-free one, omits it.

- Loss-free nominal point estimate: **24.3 kWh/kg LOX**, matching Leger's central value and inside his 1-sigma interval [18.5, 30.1].
- The loss-free Monte-Carlo distribution is centered lower (median ~20.7, mean ~21.1): Leger's value sits near our 85th percentile, not our center. We therefore report agreement as substantial **interval overlap** between two independent estimates, not as a point match; the nominal coincidence should not be over-read.
- Stage shares match: heating + reaction ~65% (Leger ~55%); water electrolysis ~30% (Leger ~38%).
- A second cross-check: Taylor and Carrier (1993) place the route at ~26 kWh/kg LOX (cross-technology range 18–35).
- The headline table charges the route the shared standing-loss term like its high-temperature siblings (nominal total 32.3 kWh/kg O_2). If continuous standing loss proves negligible at scale, the headline reverts toward the validated loss-free figure; the paired ranking already integrates over that range.

3.2 MRE against Carr 1963 and terrestrial molten-oxide electrolysis

MRE began as a pure first-principles estimate (nominal 26.5, 90% CI [20.1, 42.0], now including the reactor-loss term). Because it has no published electrical figure on this boundary, we cross-check it against three sources that unavoidably use *different system boundaries*, so the agreement is weak and bounding rather than a clean anchor: Carr (1963), as reported secondarily by Schreiner et al. (2016), gives ~26.4 kWh/kg O_2 for lunar MRE; terrestrial molten-oxide electrolysis of iron runs ∼3.7–4.0 MWh/t metal, i.e. ∼9 kWh/kg O_2, a well-insulated lower bound for a different chemistry (Allanore, 2015); and NASA reactor-sizing models place *whole-system* figures (including Joule heating and duty cycle, which our standalone-reactor boundary excludes) at ∼50–120 kWh/kg O_2 (Schreiner et al., 2016). Our estimate falls between these, which given their ∼9–120 spread is corroboration only in a loose sense. We did, however, correct the oxide-melt current efficiency to match *measured* regolith-surrogate data (70–90%, Ir anodes Allanore, 2015) rather than an optimistic assumption, a change that lowered MRE's estimate against the prior narrative.

These are the project's falsifiable tests (`tests/test_validation.py`). Only hydrogen reduction has a like-for-like published electrical figure; the others are first-principles estimates reported with wide intervals and as probabilities, not point claims.

4 Results

Table 1: Electrical-equivalent energy per route (20,000 trials).

Route	Yields	kWh/kg O_2 (nominal)	90% CI	kWh/kg propellant
PSR water mining	LOX+LH_2	15.8	13.7–20.2	**14.1**
Carbothermal (CH_4)	LOX	21.1	16.1–33.1	21.1
Molten-salt (FFC)	LOX	23.3	17.9–36.2	23.3
Molten regolith electrolysis	LOX	26.5	20.1–42.0	26.5
H_2 reduction (ilmenite)	LOX	32.3	22.0–41.9	32.3

The single value is the deterministic nominal (all parameters at their cited value); for the right-skewed routes the Monte-Carlo median can differ by a few kWh/kg, so for plant sizing prefer the median and the interval. All five routes now carry a continuous standing-loss term (Sections 2.5–2.6): the four high-temperature routes the shared reactor loss, the water route its smaller PSR analog. The high-temperature intervals are wide and strongly overlapping; they are not cleanly separable from one another. The water route sits below them, by a narrower margin than when it was exempt from capture, excavation, and standing-loss charges.

Table 2: Paired-Monte-Carlo dominance.

Route	P(cheapest)	P(worst)
PSR water mining	0.89	0.00
Carbothermal (CH_4)	0.10	0.00
Molten-salt (FFC)	0.01	0.03
Molten regolith electrolysis	0.00	0.43
H_2 reduction (ilmenite)	0.00	0.54

The water route is the cheapest in 89% of paired trials and is essentially never the worst; carbothermal takes most of the remaining 10%. Charging hydrogen reduction the shared standing loss moves it from mid-pack to the most likely worst (0.54, with MRE at 0.43): its low per-batch oxygen yield (1.4–4.4 wt%) means it processes far more hot regolith per kg O_2 than any other route, so the standing-loss and recuperation penalties hit it hardest. The high-temperature routes remain mutually inseparable within their shared uncertainty.

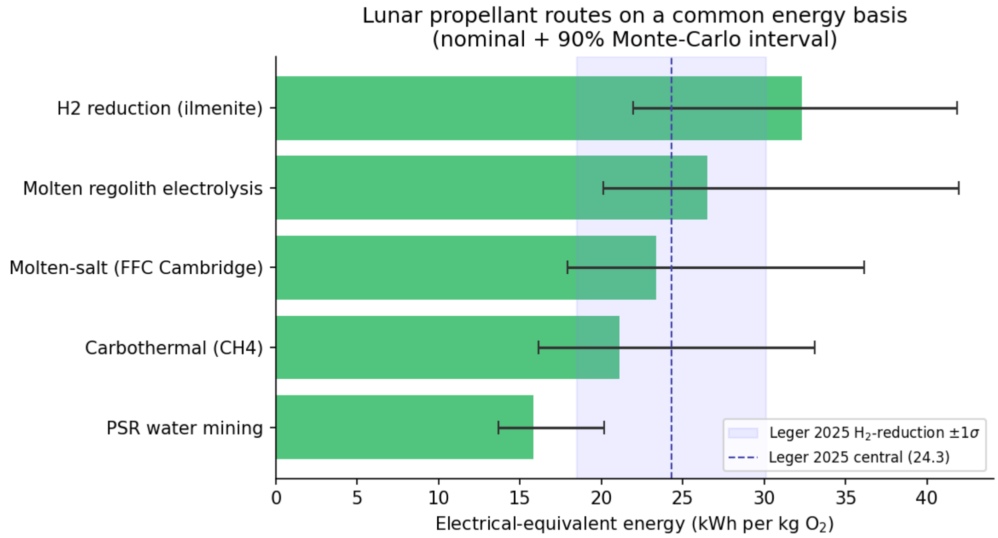

Figure 1: The five routes on a common electrical-energy basis (nominal + 90% Monte-Carlo interval); the shaded band is Leger 2025's 1-sigma for hydrogen reduction (which our loss-free configuration matches; the plotted headline charges standing loss).

5 Sensitivity analysis

A one-at-a-time tornado analysis (each parameter swept low-to-high, all others nominal; `python -m lpem -sensitivity <route>`) identifies the dominant uncertainty for each route and, with it, the single highest-value measurement:

Table 3: Dominant sensitivity driver per route.

Route	Dominant driver (swing, kWh/kg O_2)	Next
PSR water	**PSR standing loss (7.7)**	LH_2 liquefaction (6.0), ice grade (4.2), capture eff (1
Carbothermal	**reactor heat loss (28.0)**	reaction enthalpy (3.5), electrolysis eff (2.4)
Molten-salt	**reactor heat loss (28.0)**	current efficiency (5.6), cell voltage (4.5)
MRE	**reactor heat loss (28.0)**	cell voltage (15.9), yield (3.8)
H_2 reduction	**reactor heat loss (28.0)**	O_2 yield (13.3), heat recuperation (9.1), cp (5.3)

Two practical conclusions. First, the standing-loss family now dominates *every* route's uncertainty, including the winner's: the reactor term for the four high-temperature routes and the new PSR term for water. There is an irony worth stating plainly: the parameter this revision added to charge the water route honestly is now that route's largest uncertainty, which makes measuring standing loss at relevant scale (hot-reactor and PSR sublimation-zone alike) the highest-value measurement for the winning route too, ahead of the LH_2 liquefaction term that previously topped its tornado. Second, the ice-grade range now extends down to ∼1 wt% (widened from the earlier 4 wt% floor to span lean CLPA/neutron estimates), and grade now propagates through capture and excavation throughput as well as mining heat; the route ranking survives it in-band (the 0.89 dominance already integrates grade down to 1 wt%). The *Centradiant cascade thermal closure* (100 kW of compute waste heat supporting ∼1.2 t/day of water) remains grade-sensitive and degrades as grade falls toward ∼1 wt%, so that risk lives in the architecture, not in this route-energy model.

Because the one-at-a-time tornado holds other parameters at nominal and ignores interactions, we corroborate it with a global, variance-based **Sobol** decomposition (`python -m lpem`

`-sobol <route>`; Saltelli/Jansen estimators, numpy-only). We report the total-effect index S_{Ti} (variance involving an input, including its interactions) as the primary measure; the first-order index S_i (variance explained alone) is high-variance when one input dominates, so we read it as indicative only. The Sobol result confirms the tornado's dominant drivers globally: for the water route, the PSR standing loss ($S_{Ti} \approx 0.44$) and LH_2 liquefaction (≈ 0.36) lead, with ice grade next (≈ 0.07); for MRE, the reactor-loss term (≈ 0.59) and the coupled cell-voltage/efficiency latent (≈ 0.40); for hydrogen reduction, the reactor-loss term (≈ 0.71) followed by O_2 yield (≈ 0.18). The high-temperature routes are close to additive (total-effect indices sum to $\sim$1.0 and per-input interaction terms are within estimator noise of zero); the water route shows a modest positive interaction on its standing-loss term ($S_{Ti} \approx 0.44$ vs $S_i \approx 0.24$), as expected for a log-range parameter multiplying a grade-dependent throughput. The robustness claim is the modest one: the dominant uncertainties are identified by both a local (tornado) and a global (Sobol) method.

5.1 Solar-thermal process heat: the baseline is a siting assumption

The all-electric baseline (Section 2.2) charges every route's process heat through resistive heating. At a sunlit site, concentrated sunlight can deliver high-grade heat directly, and the carbothermal literature (including the CaRD program itself) assumes exactly that. Re-evaluating the routes with all high-grade thermal demand (sensible, fusion, reaction heat, and standing loss) supplied solar-thermally (`python -m lpem -solar-thermal`; concentrator mass and tracking not modeled):

Table 4: All-electric vs solar-thermal process heat (nominal kWh/kg O_2).

Route	All-electric (nominal)	Solar-thermal at sunlit site
PSR water mining	15.8	15.8 (in permanent shadow; cannot use concentrators)
Carbothermal (CH_4)	21.1	**8.2**
Molten-salt (FFC)	23.3	14.5
Molten regolith electrolysis	26.5	15.8
H_2 reduction (ilmenite)	32.3	**8.2**

The ranking inverts: with free high-grade heat, carbothermal and hydrogen reduction fall to $\sim$8 kWh-electric/kg O_2, well below the water route, whose remaining demand is overwhelmingly electrical (electrolysis, liquefaction) and whose site precludes concentrators. The honest statement of this paper's central result is therefore conditional: **the low-temperature advantage holds where process heat must be electrical** (PSR siting, nuclear-powered architectures, eclipse-tolerant continuous operation), and a sunlit solar-thermal architecture is the strongest challenger to the water route, at the cost of concentrator mass, sun tracking, thermal cycling through the lunar night, and being hundreds of kilometers from the water resource it would need for hydrogen.

6 Findings

Finding 1: On an all-electric basis, the low-temperature polar water route is the most energy-efficient AND the only full-propellant route, and it survives being charged symmetrically. It is the cheapest in 89% of paired trials (15.8 kWh/kg O_2, 14.1 per kg of propellant) and is essentially never the worst, and this now holds *after* charging it for vapor-capture efficiency, cryogenic excavation, PSR standing loss, and the ice's full thermal chain (Section 2.6), which its earlier 0.98 dominance did not. The mechanism is structural: water mining operates at the $\sim$273 K sublimation point, so its standing loss is a T^4 factor smaller than any hot reactor's and it needs no reduction chemistry at all. It is also the only route that co-produces usable hydrogen. That completeness advantage should not be oversold: LH_2 is only ∼15% of LOX+LH_2 propellant mass, so a LOX-only route with Earth-shipped hydrogen captures most of the mass leverage; completeness is a real but second-order advantage, and it is bought with the route's own dominant cost terms (small-scale LH_2 liquefaction and, outside this boundary, months-long zero-boil-off storage at 20 K).

Finding 2: Every route carries a large, scale-uncertain continuous standing loss, and symmetric charging moves the ranking. The four high-temperature routes carry the reactor-loss term (the only measured datum, NASA CaRD at ∼63–93 kWh-thermal/kg O_2 for the reduction step, is loss-dominated at demonstrated scale); the water route carries a smaller 273 K analog. Two consequences of charging it symmetrically: hydrogen reduction, previously exempted by an inconsistent calibration argument, moves from mid-pack (24.3) to nominally the most expensive route (32.3) and the most likely worst (P 0.54), because its low per-batch yield multiplies every per-kg-regolith penalty; and the water route's margin narrows (from 7 to ∼5 kWh/kg nominal against carbothermal) while its dominance drops from 0.98 to 0.89. An early version of this model omitted the term entirely and found carbothermal cheapest; a later version charged it to only three routes and found water winning at 0.98. Both were artifacts of asymmetric charging. The 0.89 figure is the defensible one.

Finding 3: The honest separation is low-temperature vs high-temperature, conditional on electric heat. Where process heat must come from electricity (permanent shadow, nuclear baseload, operation through the lunar night), the dominant energy distinction across all five routes is operating temperature, not extraction chemistry: one low-temperature route (∼16) sits below a cluster of four high-temperature routes (∼21–32) whose differences are smaller than their shared reactor-loss uncertainty. At a sunlit site with solar-thermal process heat the separation inverts (Section 5): carbothermal and hydrogen reduction drop to ∼8 kWh-electric/kg O_2 and undercut water. For an architect the first-order levers are therefore siting and heat supply, and only then chemistry: in shadow, avoid high temperature or drive down its standing loss; in sunlight, concentrated solar heat is worth more than any choice among reduction chemistries.

7 From energy to power plant and landed mass

Energy matters through what it costs to supply. Converting kWh/kg into continuous surface power and the landed mass of the fission-surface-power (FSP) system it implies (`lpem.arch`), even a modest 50 t O_2/yr plant requires ∼100–205 kWe, one to two units of NASA's planned 100-kWe FY2030 reactor (the cited 40-kWe concept (Oleson et al., 2022) scaled up). Route choice swings landed power-system mass by ∼24 t for the same oxygen output: the low-temperature water route needs ∼23 t of FSP, while the high-temperature routes need ∼30–46 t (carbothermal ∼30, molten-salt ∼33, MRE ∼38, hydrogen reduction ∼46), the spread driven by their reactor-loss burden and, for hydrogen reduction, its yield-multiplied throughput. The direction is robust; the magnitude scales with the FSP specific mass and assumes an all-fission architecture (a solar-plus-storage plant would be dominated by night-survival storage mass, not modeled here; a solar-thermal process-heat architecture would change the comparison per Section 5).

8 Companion analyses (compute waste-heat integration)

Two companion analyses, kept modular, examine reusing co-located compute waste heat:

- **Low-grade thermal offset** (`WASTE-HEAT-OFFSET.md`). By the Second Law, compute waste heat (∼330 K) can supply only low-grade ISRU demand. The offset model now derates delivery through a heat-exchanger approach ΔT (nominal 15 K) and effectiveness (nominal 0.85). At the nominal 350 K reject temperature, comfortably above the ∼273

K sublimation target even after the pinch, it can cover the water route's thermal-mining term (∼1.8 kWh/kg O_2, ∼11% of the route; Figure 2); this offset falls off sharply and excludes the sublimation enthalpy if the post-pinch source temperature drops below ∼273 K. It cannot supply high-grade reduction/electrolysis heat at all, and the supportable-throughput figures it implies are energy-balance upper bounds: delivering the heat into a granular icy bed is conduction-limited, which the model states but does not solve.

- **PSR co-location** (`PSR-COLOCATION.md`). A permanently shadowed region is both a cryogenic radiative sink for compute and the location of the water resource. Under an explicit radiator energy balance with a sky view factor (a competently oriented vertical two-sided panel, F_{sky} nominal 0.5, replacing an earlier horizontal-panel assumption that overstated the case), PSR siting saves a feasible-case median ∼10 t of radiator mass per MW of compute (IQR ∼7–16; Figure 3); only ∼0.1% of sampled sunlit designs truly cannot reject at 330 K (a further ∼0.4% need a prohibitive >10x area), correcting an earlier claim of ∼30% that was an artifact of the horizontal-panel geometry. The heat cascade itself is a modest bonus (saves ∼2.3 t reactor for the reference plant against ∼1 t of integration hardware; the break-even enabling probability is ∼44% at nominal, ∼50% propagated, with P(worthwhile once co-located) ∼78%); co-location is justified by the compute siting economics, not the cascade. This is, fittingly, most useful at the low-temperature water route that Finding 1 already favors.

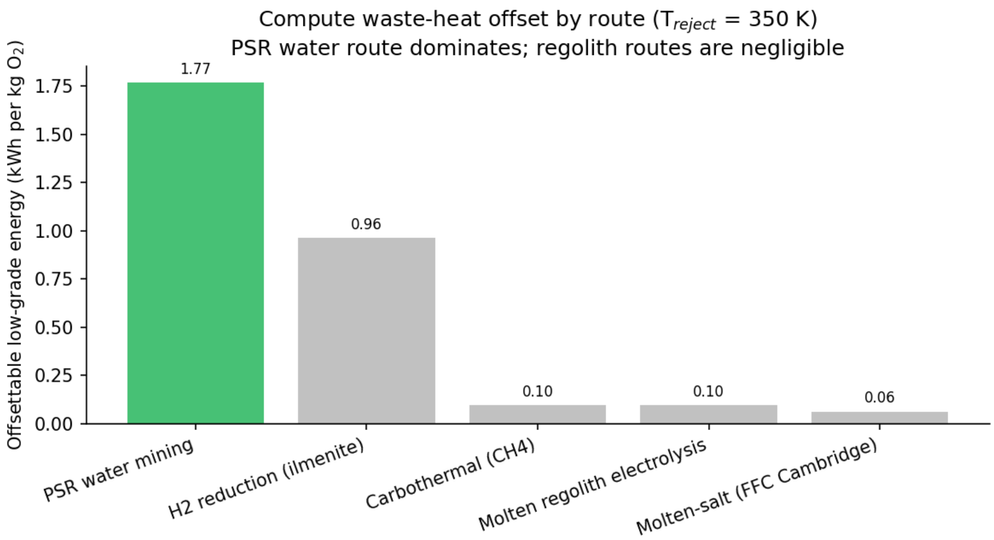

Figure 2: Low-grade, waste-heat-offsettable energy per route ($T_{\mathrm{reject}} = 350$ K, pinch and effectiveness applied).

9 Limitations

- Steady-state production energy, power, and FSP landed mass only; capital cost, mobility, comms, site logistics, and demand are out of scope.
- The largest previously-omitted water-route terms are now charged (capture efficiency, cryogenic excavation, PSR standing loss, ice thermal chain; Section 2.6). Terms still omitted are one-signed (they raise totals) and still fall disproportionately on the water route: months-long LH_2 zero-boil-off storage, power delivery into permanent shadow, mobility and haulage in a PSR, and survival heating for equipment at 40–110 K. Omissions

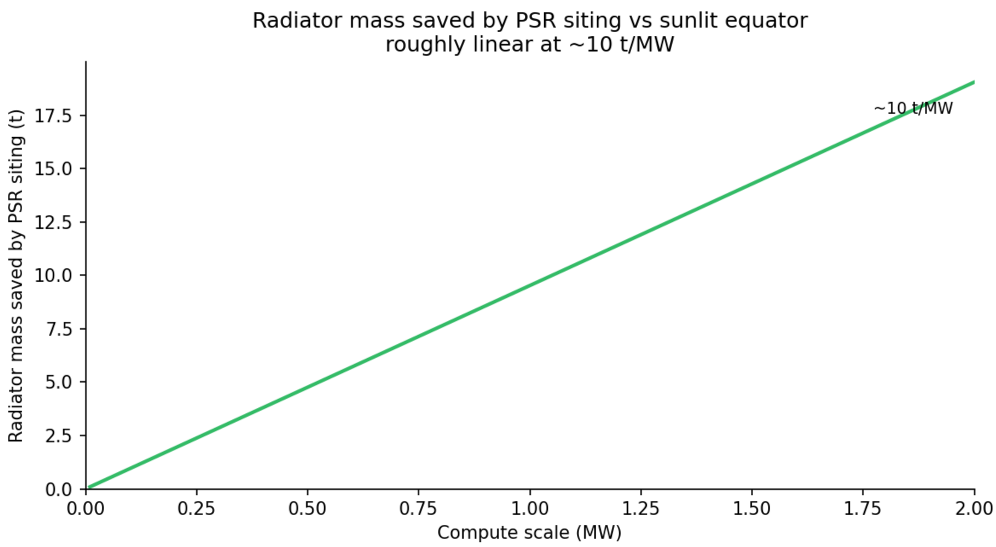

Figure 3: Radiator mass saved by PSR siting vs a sunlit vertical-panel site (feasible-case median ∼10 t/MW; the saving is right-skewed).

against the other routes: comminution/beneficiation and electrolyzer heat-rejection parasitics. Absolute kWh/kg figures are best read as lower bounds; the ranking is more robust than the levels, but the remaining asymmetry means the water route's 0.89 dominance should be read as an upper bound on confidence, not a lower one.

- The solar-thermal sensitivity (Section 5) omits concentrator mass, pointing, and night-survival thermal cycling; it brackets the sunlit case rather than designing it.
- Three of five routes lack a direct independent electrical anchor; the winning route's only anchor is this model itself.
- O_2 yield and reaction temperature are sampled independently (a residual idealization).
- Site geography is not modeled beyond the shadow/sunlit distinction.

10 Reproducibility

```
pip install -e .
python -m lpem                  # Table 1
python -m lpem --dominance      # Table 2
python -m lpem --sensitivity mre    # Section 5 tornado
python -m lpem --sobol mre           # Section 5 Sobol variance decomposition
python -m lpem --solar-thermal       # Section 5 solar-thermal sensitivity
python -m lpem --plant-tonnes 50    # Section 7
python -m lpem --waste-heat --benefit    # Section 8
python scripts/make_figures.py      # Figures 1-3
pytest                          # 68 tests, incl. the validation anchors
```

References

See `paper/REFERENCES.md` (full bibliography with DOIs/URLs) and `paper/references.bib`. Key sources: Leger et al. (2025) (PNAS); Taylor and Carrier (1993); Carr (1963); Allanore

(2015) (J. Electrochem. Soc.); Schreiner et al. (2016) (Adv. Space Res.); Kornuta et al. (2019) (REACH); Hauch et al. (2020) (Science); Colaprete et al. (2010) (Science).

References

Antoine Allanore. Features and challenges of molten oxide electrolytes for metal extraction. *Journal of the Electrochemical Society*, 162(1):E13–E22, 2015. doi: 10.1149/2.0451501jes. Terrestrial MOE voltages and current efficiency; 3.7 MWh/t Fe (9 kWh/kg O2) lower bound for MRE.

B. B. Carr. Recovery of water or oxygen by reduction of lunar rock. *AIAA Journal*, 1(4):921–924, 1963. doi: 10.2514/3.1674. Earliest lunar oxygen energy estimate. The 26.4 kWh/kg O2 figure is reported by secondary sources (Schreiner 2016), not lifted verbatim from this paper.

Anthony Colaprete, Peter Schultz, Jennifer Heldmann, Diane Wooden, Mark Shirley, Kimberly Ennico, et al. Detection of water in the LCROSS ejecta plume. *Science*, 330(6003):463–468, 2010. doi: 10.1126/science.1186986. LCROSS PSR ice grade 5.6 +/- 2.9 wtHeld as raw/colaprete-2010-lcross-water-detection.pdf.

Anne Hauch, Rainer Küngas, Peter Blennow, Anders B. Hansen, John B. Hansen, Brian V. Mathiesen, and Mogens B. Mogensen. Recent advances in solid oxide cell technology for electrolysis. *Science*, 370(6513):eaba6118, 2020. doi: 10.1126/science.aba6118. SOEC system efficiency (0.67); water-electrolysis efficiency parameter, set independently of Leger.

International Energy Agency. The future of hydrogen: Seizing today's opportunities. Technical report, IEA, Paris, 2019. URL https://www.iea.org/reports/the-future-of-hydrogen. Report prepared for the G20, Japan. Electrolysis system-efficiency benchmark for the water-electrolysis stage.

David Kornuta, Angel Abbud-Madrid, Jared Atkinson, Justin Barr, Gary Barnhard, Dallas Bienhoff, et al. Commercial lunar propellant architecture: A collaborative study of lunar propellant production. *REACH – Reviews in Human Space Exploration*, 13:100026, 2019. doi: 10.1016/j.reach.2019.100026. CLPA: 450 t/yr water plant at 2.0 MWe; 2.8 MW plant. Cross-check for power/landed-mass module. Held as raw/clpa-kornuta-2019.pdf.

Dorian Leger, Freja Thoresen, Fardin Ghaffari-Tabrizi, Matthew Shaw, Joshua Rasera, David Dickson, Baptiste Valentin, Anton Morlock, and Aidan Cowley. Modeling energy requirements for oxygen production on the Moon. *Proceedings of the National Academy of Sciences*, 122 (8):e2306146122, 2025. doi: 10.1073/pnas.2306146122. Primary anchor: H2 reduction 24.3 +/- 5.8 kWh/kg LOX; stage breakdown; 11.3 kWh/kg O2 water benchmark. Verified against raw/leger-2025-oxygen-production-energy-model-pnas.pdf.

Steven Oleson, Thomas Packard, Elizabeth Turnbull, Marc Gibson, Dasari Rao, Christopher Barth, Scott Wilson, Paul Schmitz, Anthony Colozza, Brandon Klefman, Lucia Tian, and Lee Mason. A deployable 40 kWe lunar fission surface power concept. In *Nuclear and Emerging Technologies for Space (NETS) 2022*. NASA Glenn Research Center (Compass Team), 2022. URL https://ntrs.nasa.gov/citations/20220004670. NASA NTRS Document ID 20220004670. FSP 40 kWe concept; specific-mass basis (225 kg/kWe; 150 kg/kWe target). DOI was an ANS placeholder in the manuscript.

Samuel S. Schreiner, Laurent Sibille, Jesus A. Dominguez, and Jeffrey A. Hoffman. A parametric sizing model for molten regolith electrolysis reactors to produce oxygen on the Moon. *Advances in Space Research*, 57(7):1585–1603, 2016. doi: 10.1016/j.asr.2016.01.029. MRE reactor sizing; whole-system kWh/kg O2 (50-120) with Joule heating and duty cycle.

Lawrence A. Taylor and III Carrier, W. David. Oxygen production on the Moon: An overview and evaluation. In John S. Lewis, Mildred S. Matthews, and Mary L. Guerrieri, editors, *Resources of Near-Earth Space*, page 69. The University of Arizona Press, 1993. Independent anchor: H2 reduction 26 kWh/kg LOX; cross-tech range 18-35 kWh/kg O2 over 20 reviewed processes. ADS bibcode 1993rnes.book...69T.