A Differentiable Multi-Fidelity Model for IFE Design

A proof of concept for vertically-integrated fusion design

Archis Joglekar

May 5, 2026

Pasteur Labs, Ergodic LLC

Acknowledgments

Collaborators

- J. Brodrick, J. Coughlin, A. Lavin (Pasteur Labs)
- A. Crilly (Imperial College London)
- P. Travis (Ergodic LLC)

Funding

- Department of Energy - Fusion Energy Sciences
- Pasteur Labs
- Schmidt Fellowship AI for Science

Starting point: GEM is an Excel-based techno-economic model

GEM (LLNL) couples target physics, driver, chamber, BOP, and economics into a single consistent framework for IFE design trades and sensitivity studies.

Released as an Excel workbook: tabs for each subsystem, parameter sweeps driven by hand or via VBA.

We re-implement GEM in JAX and make the entire chain differentiable end-to-end.

Inputs	DISCLAIMER: This software is intended fo do not reflect assumptions or projections		
Primary Design Variables	**User Inputs**	**Units**	**Notes**
Laser energy on target	5.00	MJ	*Value must be >*
Pulse repetition frequency	10.00	Hz	*Resets to last ca*
Optional Plant Size Variable			
Net electric power to grid	1000.0	MWe	*Only used when*
	Fix Net power		*Resets pulse rep*
Operating and Cost Multipliers			
Gain curve multiplier	0.25	-	
Driver efficiency multiplier	0.85	-	
Driver cost multiplier	1.00	-	
Target cost multiplier	1.00	-	
Secondary Design Variables			
Plant type selection	2	-	*1 = LIFE, 2 = HAI*
Thermal conversion efficiency	40.0%	-	*Efficiency of co*
Auxiliary power consumption	5.0%	-	*Portion of gross*
Plant availability	85.0%	-	*Portion of year p*
Economic Design Variables			
Cost evaluation year	2025	-	*Cost estimates*

The question

What additional information can we extract from this cost-accounting framework if we make the entire chain differentiable?

- The full Jacobian $\partial(\mathrm{COE})/\partial\boldsymbol{\theta}$ across all parameters simultaneously
- Uncertainty attribution through the Jacobian and an assumed covariance over the input parameters
- Gradient-based optimization over the coupled design space

What this talk is — and isn't

This is a proof of concept. The contribution is the *pattern*, not the numbers — one differentiable graph across the full IFE stack.

Caveats:

- I'm a PDE-solver person - past 4 years = differentiable Vlasov x 2, differentiable Thomson scattering, differentiable laser propagation — *not* a fusion economist, so I defer to the literature for cost data, learning rates, TRL anchors, etc.
- Cost data, learning rates, TRL anchors all need refresh — the framework is indifferent to which numbers you put in
- A full-stack fusion org plugs in its own driver, target physics, BOP — the API guarantees between them hold

Claiming: methodology and leverage. **Not claiming:** 245.6 \$/MWh is the answer.

Why this matters now

Economic context:

- LCOE target for competitiveness: \$80–100/MWh (Lindley 2023, Schwartz 2023)
- Global system value of fusion: \$3.6–8.7 trillion (MIT Energy Initiative 2024)
- $>$50% value loss from 15-year delay (Spitzer et al. 2025)
- Learning rates likely 2–8%, not 8–20% (Tang et al. 2026, Nature Energy)

Implication:

Low learning rates mean cost reduction through fleet deployment is limited.

End-to-end design optimization becomes more important.

Efficient exploration of the coupled design space is valuable.

Two reference operating points: LIFE and HAPL

We use two design points throughout the talk. They differ in driver, chamber, and assumed gain.

	LIFE LLNL, 2009–13	**HAPL** NRL-led, 2000s
Drive	Indirect (hohlraum)	Direct
Chamber	Wetted-wall (Pb-Li)	Dry-wall
Laser energy	1.85 MJ	5 MJ
Rep rate	10 Hz	10 Hz
Gain at design	≈ 39	≈ 52
m (gain mult.)	0.846	0.25
Thermal eff.	0.419	0.40

Gain curve: the Amendt 2011 power law $G(E) = m \cdot 60 \cdot (E/2.2\text{ MJ})^{1.505}$, scaled by m.

Provenance:

- **HAPL:** GEM v1.0 spreadsheet defaults (LLNL-CODE-2013481). $m = 0.25$: a conservative anchor — $\frac{1}{4}$ of the (indirect-drive) Amendt curve, since direct-drive ignition at MJ scale is unproven.
- **LIFE:** *calibrated by us* to LIFE.2 (gain ≈ 39 at 1.85 MJ); back-solve $m = 0.846$ so the Amendt curve passes through that point.

The Framework

Differentiable GEM: decomposed into a DAG of Differentiable Code Artifacts

GEM (LLNL): IFE techno-economic model — target physics, laser driver, chamber + blanket + tritium, BOP, capital rollup to COE.

We compose it as a DAG of nodes at discipline boundaries (driver ∥ nuclear island), so each block can be swapped or replaced with a higher- or lower-fidelity model

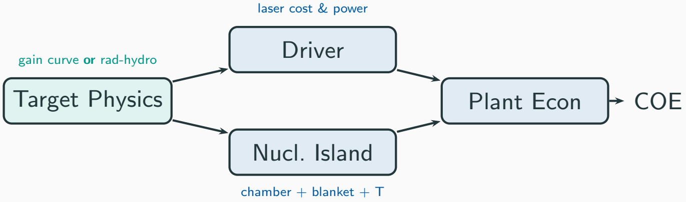

Tesseracts: a contract for composable simulation

The problem a system-of-systems model has to solve: target physics, driver, chamber, and BOP are owned by different teams, written in different stacks, and evolve at different rates.

A Tesseract is a component with a fixed contract:

- `InputSchema` / `OutputSchema` (Pydantic — typed boundaries)
- `apply(inputs)` — forward
- `vjp(inputs, cotangent)` — reverse-mode derivative
- `jacobian(inputs)` — full Jacobian

What you get for free:

- **Interoperability:** schemas are the API; teams ship independently
- **Fidelity swapping:** replace a scaling law with a PDE solver — same contract, gradients still flow through
- **Composition:** chain rule *is* the system integrator — no glue code per pair

This talk: GEM decomposed into four Tesseracts (Target Physics → Driver ∥ Nuclear Island → Plant Economics).

`github.com/pasteurlabs/tesseract`

Sensitivity ranking

Full sensitivity ranking

One backward pass gives the elasticity of COE to all 15 design parameters simultaneously.

	Parameter	Elasticity
1.	Laser energy	-1.88
2.	Thermal efficiency	-1.45
3.	Gain curve multiplier	-1.20
4.	Fixed charge rate	$+0.74$
5.	Rep rate	-0.64
6.	Driver efficiency	-0.58
7.	Driver cost	$+0.52$
	Remaining 8 parameters: $\lvert\varepsilon\rvert < 0.18$ (O&M, contingency, target cost, TBR, ...)	

Elasticity $\varepsilon = (\partial \mathrm{COE}/\partial\theta) \cdot (\theta/\mathrm{COE})$. Target physics (laser energy, gain multiplier) and the thermal cycle dominate; driver-side parameters and the fixed charge rate are second-tier; the remaining 8 are below 0.18.

Where does model fidelity matter?

Gradient magnitudes show where higher-fidelity models pay off. *Conditional on LLNL cost data and $m = 0.25$.*

High elasticity (fidelity matters):

- Laser energy: $|\varepsilon| = 1.88$
- Thermal efficiency: $|\varepsilon| = 1.45$
- Gain multiplier: $|\varepsilon| = 1.20$

Low elasticity (scaling laws adequate):

- TBR: $|\varepsilon| = 0.007$
- Target fab cost: $|\varepsilon| = 0.08$
- Auxiliary power: $|\varepsilon| = 0.08$

The case for rad-hydro:

The empirical gain curve maps laser energy to gain — one input. A rad-hydro solver maps the full *pulse shape*: temporal profile, picket/foot power, rise time.

Target physics has high elasticity, so this expanded space is worth exploring.

Uncertainty attribution

Uncertainty attribution — which speculative costs matter?

Since the Jacobian is already computed, first-order variance attribution comes for free:

$$\sigma^2_{\mathrm{COE}} \approx J\,\Sigma_\theta\,J^\top \qquad \Longrightarrow \qquad f_i = \frac{(\partial \mathrm{COE}/\partial\theta_i)^2\,\sigma_i^2}{\sigma^2_{\mathrm{COE}}}$$

Assumed σ_i ($\pm 1\sigma$ relative to nominal):

Parameter	σ_i/θ_i	**Source**
Gain multiplier	$\pm 50\%$	Woodruff 2026 (factor $\sim$2 at low TRL)
Driver cost mult.	$\pm 50\%$	Hawker 2020 (log-normal $\sigma \approx 0.5$)
Driver eff. mult.	$\pm 30\%$	Meier 2009 (0.5–2$\times$ swept); diode TRL 4–5
Thermal eff.	$\pm 10\%$	GEM v1.0 user guide (0.35–0.45)
Fixed charge rate	$\pm 20\%$	GEM v1.0 user guide (0.06–0.09)

Uncertainty attribution: the ranking shifts with operating point

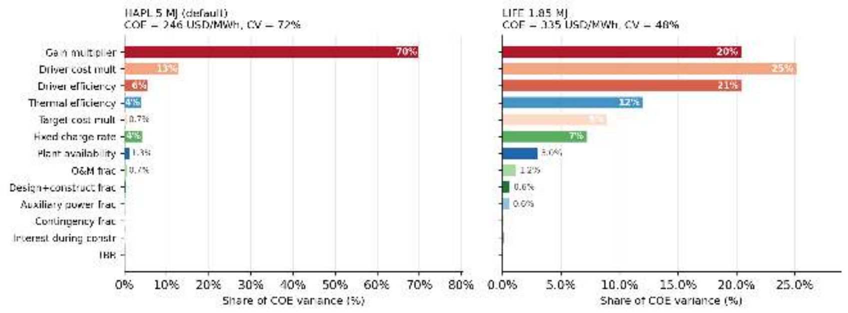

HAPL 5 MJ (low gain): gain curve uncertainty alone drives ∼70% of COE variance — reducing it has the highest payoff. **LIFE 1.85 MJ** (high gain): driver cost (∼25%), driver efficiency (∼21%), and gain multiplier (∼20%) contribute roughly equally — the bottleneck shifts to driver technology.

Optimization & robust design

Gradient-based plant optimization

L-BFGS over E_{laser} at fixed 1000 MWe — rep rate falls out of the power constraint.
($\eta_{\text{driver}}, \eta_{\text{thermal}}$, TBR are monotonic; deferred to robust-design slide where uncertainty turns them into real tradeoffs.)

	LIFE ref / opt	HAPL ref / opt
E (MJ)	1.85 / 2.55	5.0 / 3.7
Rep rate (Hz)	51 / 11	15 / 18
Recirc	40% / 8%	35% / 17%
Capital ($B)	11.0 / 7.4	14.0 / 9.7
COE ($/MWh)	195 / **107**	194 / **144**

COE: LIFE −45%, HAPL −26%

One real tradeoff: E vs. rep rate.

- Raising E raises driver capital — pull *up*
- Higher E $\rightarrow$ higher gain $\rightarrow$ rep rate drops to hold 1000 MWe $\rightarrow$ less recirculation — pull *down*
- Optimizer lands where the two pulls balance. Next slide: the gradient decomposition that proves it's an equilibrium.

Why does the optimizer move there?

$d(\mathrm{COE})/dE$ splits into three channels — the two pulls from before, plus a cross-term where rep rate does double duty:

HAPL reference ($E = 5$ MJ):

Channel	d**COE**$/dE$
Driver cost ($E \to$ cost)	$+5.9$
Yield/BOP ($E \to G \to Y$)	-15.6
Cross-terms (rep rate coupling)	$+11.3$
Total	$+1.5$

$\rightarrow$ Yield path 2.7$\times$ stronger than driver path; reference is near-optimal.

At the optimum:

Channel	d**COE**$/dE$
Driver cost	$+5.3$
Yield/BOP	-17.1
Cross-terms	$+11.9$
Total	≈ 0

$\rightarrow$ Driver-cost increase and yield benefit exactly cancel: an equilibrium.

Robust design: technology targets aren't realized values

Setup. Realized efficiency carries uncertainty that grows as you target performance past demonstrated SOTA.
Illustrative functional form: $\sigma_\eta(\eta_{\text{design}}) = \sigma_{\text{floor}} + \alpha \cdot \max(0, \eta_{\text{design}} - \eta_{\text{SOTA}})^3$.
Mean–variance objective: $F_\lambda(x) = \mathbb{E}[\text{COE}](x) + \lambda\, \sigma_{\text{COE}}(x)$ with $\sigma^2_{\text{COE}} = J^\top \Sigma(x) J$, $J = \nabla_\theta \text{COE}$ — needs $\nabla_x \sigma_{\text{COE}}$, a derivative *through* a derivative.

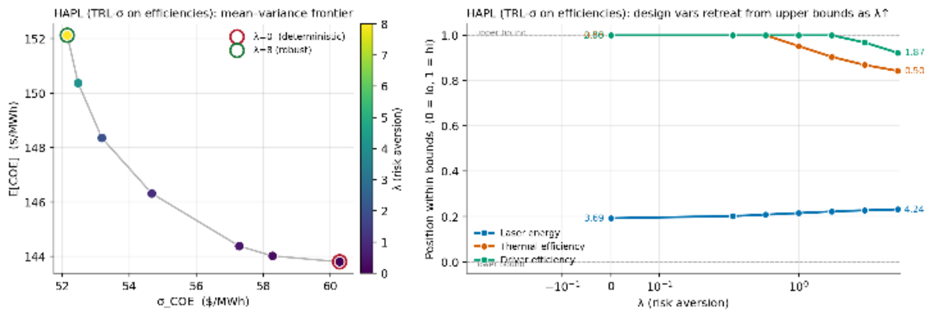

Where does the variance reduction come from?

Cost of robustness (HAPL)

	λ=0	λ=1	λ=2	λ=8
$\mathbb{E}$[COE] ($/MWh)	144	146	148	**152**
σ_{COE} ($/MWh)	60	55	53	**52**

At λ=8: +$8/MWh in expected COE buys −13.5% in σ. LIFE buys −26.5% for +$11/MWh.

Variance attribution flips (HAPL share of σ^2)

	λ=0	λ=8	
Thermal eff. (TRL-σ)	30%	7%	▼
Driver eff. (TRL-σ)	10%	6%	▼
Driver cost (exogenous)	28%	49%	▲
Fixed charge rate (exogenous)	11%	17%	▲
Gain mult. (exogenous)	8%	8%	

Robust optimum is *interior*: efficiencies retreat from the bounds until the marginal expected-cost of retreating equals the marginal σ-reduction it buys.

The residual σ is dominated by **exogenous** cost uncertainties that no design choice can reduce → that's the irreducible risk floor.

Inverse design

Inverse design — min-norm bundle of breakthroughs

New question. Anchor against *demonstrated* SOTA (not LIFE/HAPL design references). For a target COE, what's the cheapest joint push from SOTA that gets there? Each capability has a different "one TRL step" size; weight excursions in those units and minimize the sum.

$$\min_{\theta \in \mathcal{D}} \sum_{i \in \mathcal{D}} w_i \left[\mathrm{ReLU}(\theta_i - \theta_{\mathrm{SOTA},i})\right]^2 \quad \text{s.t.} \quad \mathrm{COE}(\theta) \leq \mathrm{COE}_{\mathrm{target}}$$

- $\mathcal{D} = \{$gain mult, η_{driver}, η_{thermal}, rep rate$\}$; E fixed at 2 or 5 MJ (facility-scale, not a TRL move)
- $w_i = 1/(\Delta_i^{\mathrm{ref}})^2$ — so $w_i(\Delta\theta_i)^2$ is each capability's share of budget in squared-TRL-steps
- rep rate SOTA is anchor-dependent: **NIF** (0.001 Hz, single-shot) vs **Mercury** (5 Hz, lab rep) — run both
- Penalty method, $\mu \in \{10^2, \ldots, 10^5\}$, L-BFGS-B via `jax.grad` through coupled physics

TRL-cost stack at three target COEs

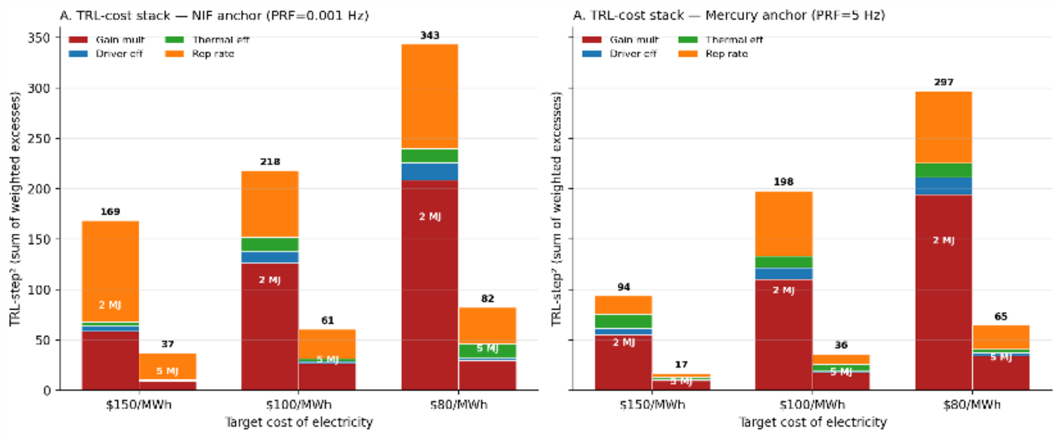

Bar height = total squared-TRL-step effort; segments = per-capability share. NIF anchor: most of the LIFE/2 MJ budget is rep rate; HAPL/5 MJ is mostly insulated. **Tightening from \$150 → \$80 roughly doubles the budget at every anchor.**

Required θ vs SOTA — NIF anchor

HAPL 5 MJ, three COE targets per capability. Bar height $= \theta_i^* / \theta_{i,\mathrm{SOTA}}$ on a log axis; absolute optimum θ_i^* printed above each bar.

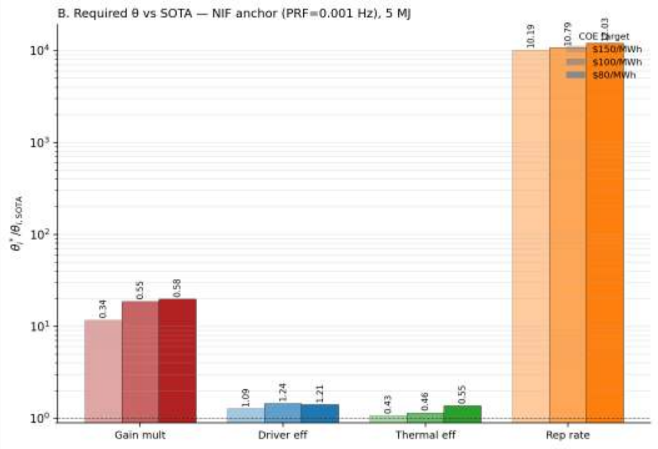

NIF SOTA = 0.001 Hz

- **Rep rate dwarfs everything** — $\sim 10^4\times$ ratio because rep rate must climb from 0.001 Hz to $\sim$10 Hz for a working plant
- Gain pushed 12–21$\times$ SOTA (fusion gain 70 $\rightarrow$ 120 at 5 MJ across the three targets)
- Driver / thermal pushes are modest — η_{driver} under 1.3$\times$, η_{thermal} saturates the bound only at $80

Required θ vs SOTA — Mercury anchor

Same axes, rep rate SOTA replaced with 5 Hz (lab demonstration). Granting lab-rep technology compresses the rep-rate bar to a $\sim$ 2–3$\times$ ratio.

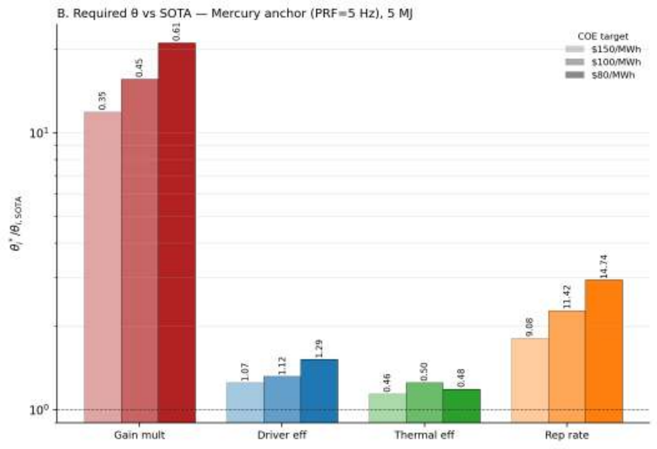

Mercury SOTA = 5 Hz

- Rep rate is now in the same magnitude class as the other capabilities — the bottleneck collapses to gain
- Gain ratios match the NIF anchor closely — the COE constraint demands the same physics regardless of where rep rate started
- **Grant rep-rate technology and the economic gap is about ICF physics, not repetition**

Composing through a PDE

lagradept: a differentiable rad-hydro PDE in JAX

A 1D Lagrangian radiation-hydrodynamics solver for ICF implosions, written end-to-end in JAX. Drops into the Target Physics node of the DAG with *no other changes* to the rest of the chain.

Physics in the loop:

- Compressible Euler hydrodynamics
- Flux-limited Spitzer thermal conduction (electron & ion)
- Single-group radiation diffusion
- Tabular FEOS equation of state (DT)
- Sector ray-traced laser deposition (Follett 2025)

What you get:

- Forward: ∼80 s for a 480k-step OMEGA-scale pulse
- Reverse: `jax.grad` flows through every step via `RecursiveCheckpointAdjoint`
- Same Tesseract contract as the empirical gain curve: same `InputSchema` / `OutputSchema` / `VJP` interface

Composes, doesn't replace — gain curve and PDE are both valid Target Physics nodes.

The parameter space expands dramatically using lagradept

Empirical gain curve

$$\text{gain}(E_{\text{laser}}) = G_0 \left(\frac{E_{\text{laser}}}{E_{\text{ref}}} \right)^{\alpha}$$

Differentiable design space:

- E_{laser} (1 scalar)

One number in, one number out.

lagradept

$$\text{gain}\left[P_{\text{laser}}(t),\ \text{beam},\ \text{target},\ \text{physics} \right]$$

Differentiable design space:

- **Pulse:** $P_{\text{laser}}(t)$ shape + peak + total energy (∼**500** knots)
- **Beam:** central wavelength, focal radius, beam count, geometry
- **Target:** r_{max}, ice thickness, gas density, ablator
- **Physics:** flux limiters, α_f, resonant-abs. frac

∼**510+ differentiable inputs**, returning gain, yield, burn fraction, ρR, burn time, σ_t.

Open-ended exploration of direct-drive driver architectures:
Excimer energy (two-sided) · Blue Laser Fusion (∼351 nm, $\mathcal{O}(100)$ beams) · ...

Pricing a hidden physics knob: the electron flux limiter

The electron flux limiter f_e is a closure parameter on Spitzer heat flow — ICF practice spans $[0.05, 0.15]$.

Sweep through HAPL 5 MJ:

f_e	COE [$/MWh]
0.05	231.0
0.10 (baseline)	245.6
0.15	262.9

±$17**/MWh (**±7%**) on COE** from a single physics-closure parameter — on par with the top engineering elasticities.

Summary

Summary

Five results from one differentiable graph, at no extra modeling cost:

1. **Sensitivity ranking:** 16 elasticities in one backward pass — laser energy (−1.88), thermal eff. (−1.45), gain mult. (−1.20) dominate.
2. **Uncertainty attribution:** $J\Sigma J^\top$ free. Gain curve drives 70% of COE variance at low gain; at high gain driver cost / efficiency / gain mult. contribute equally.
3. **Optimization & robust design:** L-BFGS cuts COE 26–45% vs. reference (1000 MWe). TRL-aware σ_η replaces "push to the bound" with an interior robust optimum; σ_{COE} down 13–27%.
4. **Inverse design (min-norm TRL push):** NIF anchor — rep rate dominates ($10^4\times$). Mercury anchor — bottleneck collapses to gain: granted lab-rep, the gap is about ICF physics.
5. **PDE-composed target physics:** `lagradept` swaps the 1-input gain curve for a ∼510-input rad-hydro solver under the same contract. Electron flux limiter alone prices at ±$17/MWh (±7%).

The numbers are placeholders; the pattern is the point. A full-stack fusion org can run this loop with its own models behind the same Tesseract contracts.

Thank you

Questions?

Backup: GEM decomposition map

Sub-tesseract	GEM Tabs	Inputs	Outputs
Target Physics	(gain curve)	laser energy, gain mult. *or* pulse + target + physics knobs	gain, yield
Driver	9	laser energy, rep rate, eval year, cost/eff mult.	driver cost, power, efficiency
Nuclear Island	6–7, 10–14	yield, laser energy, rep rate, plant type, TBR	capital, power, chamber radius, daily T
Plant Economics	(rollup)	driver + NI outputs, econ params	COE, net power, total capital

Driver and Nuclear Island have no dependencies on each other — they run in parallel.

Monolithic and decomposed DAGs produce identical COE and gradients (verified to machine precision).

Backup: Empirical gain curve scan

E_{laser} (MJ)	Gain	Yield (MJ)	COE ($/MWh)	P_{net} (MWe)
3.0	23.9	71.8	1391.8	62.0
4.0	36.9	147.5	409.0	297.4
5.0	51.6	258.0	245.6	676.6
6.0	67.9	407.4	182.4	1216.9
7.0	85.6	599.4	150.5	1933.8
8.0	104.7	837.6	132.3	2841.7

COE monotonically decreasing with E_{laser} through the full coupled system.

Amendt 2011 power-law, gain_curve_multiplier = 0.25. HAPL plant type, 10 Hz, 2025 eval year.

Backup: Sensitivities shift across the design space

Elasticities at six operating points. Some parameters are stable; others vary substantially:

Parameter	5 MJ	3 MJ	8 MJ	20 Hz	LIFE 5MJ	1.85 MJ	
Laser energy	−1.88	−7.13	−0.85	−1.78	−1.78	−2.40	▲
Thermal eff.	−1.45	−4.72	−1.03	−1.38	−1.51	−1.73	▲
Gain multiplier	−1.20	−4.61	−0.58	−1.11	−1.13	−1.46	▲
Plant avail.	−1.00	−1.00	−1.00	−1.00	−1.00	−1.00	
Fixed charge rate	+0.74	+0.68	+0.78	+0.72	+0.74	+0.64	
Rep rate	−0.64	−0.70	−0.58	−0.48	−0.52	−0.52	
Driver efficiency	−0.58	**−3.80**	−0.22	−0.58	−0.51	−0.73	▲
Driver cost	+0.52	+0.60	+0.36	+0.38	+0.50	+0.48	

▲ = varies >3× across points. First four columns: HAPL plant type. "LIFE 5MJ": LIFE plant type at 5 MJ. "1.85 MJ": LIFE plant type at 1.85 MJ, 10 Hz (low-energy regime).

Example: driver efficiency elasticity is −3.80 at 3 MJ but −0.22 at 8 MJ (17×).

Backup: Local gradient vs. finite-difference sweep

Sweeping 5 parameters over $[0.5\times, 2\times]$ gives the average COE response. How does this compare with the local gradient at the reference point?

Parameter	**COE @0.5×**	**COE @2×**	**AD/FD**	
Rep rate	403.8	166.5	1.00	linear
Driver cost	182.2	372.3	1.00	linear
Target cost	235.3	266.2	1.00	linear
Driver efficiency	584.1	190.4	**0.54**	nonlinear
Thermal efficiency	1090	107.9	**0.54**	nonlinear

AD/FD $= 0.54$: the local gradient is about half the average slope over the sweep range. This indicates strong curvature — the COE response to these parameters is nonlinear, so the marginal sensitivity at the reference point differs from the average over $[0.5\times, 2\times]$. Both numbers are correct; they answer different questions.

Backup: Full elasticity table at 5 MJ reference

	Parameter	Value	$d\text{COE}/d\theta$	Elasticity
1.	Laser energy (MJ)	5.0	-92.3	-1.88
2.	Thermal efficiency	0.40	-887.6	-1.45
3.	Gain multiplier	0.25	-1180	-1.20
4.	Plant availability	0.85	-288.9	-1.00
5.	Fixed charge rate	0.075	$+2420$	$+0.74$
6.	Rep rate (Hz)	10	-15.8	-0.64
7.	Driver efficiency mult	0.85	-167.4	-0.58
8.	Driver cost mult	1.0	$+126.7$	$+0.52$
9.	O&M frac	0.03	$+1449$	$+0.18$
10.	Design+construct frac	0.35	$+108.7$	$+0.15$
11.	Plant type (HAPL)	2.0	$+14.9$	$+0.12$
12.	Contingency frac	0.20	$+108.7$	$+0.09$
13.	Target cost mult	1.0	$+20.6$	$+0.08$
14.	Auxiliary power frac	0.05	$+408.3$	$+0.08$
15.	Interest during constr	0.12	$+108.7$	$+0.05$
16.	TBR	1.1	$+1.6$	$+0.01$

Backup: Minimum gain for target COE

What fusion gain is needed at each operating point? LIFE plant for $E \leq 3$ MJ, HAPL for $E > 3$ MJ

	Minimum fusion gain			
Laser energy	**5 Hz**	**10 Hz**	**15 Hz**	**20 Hz**
Target: \$150/MWh				
2 MJ	122	81	68	61
5 MJ	194	87	66	57
8 MJ	239	86	65	55
Target: \$100/MWh				
2 MJ	234	150	122	108
5 MJ	—	189	119	95
8 MJ	—	215	119	93
Target: \$80/MWh				
2 MJ	395	249	201	176
5 MJ	—	—	215	152

Backup: feasibility contours in (gain, rep rate) space

Pre-min-norm view: fix driver / thermal / TBR; sweep gain and rep rate; level-sets of COE.

HAPL plant, $E = 5$ MJ

- At 5 MJ, 10 Hz: need gain > 87 for \$150/MWh; gain > 189 for \$100/MWh
- NIF ignition (gain ≈ 1.5) is $\sim 60\times$ below the viability threshold
- LIFE target (gain 39) still falls short at most rep rates

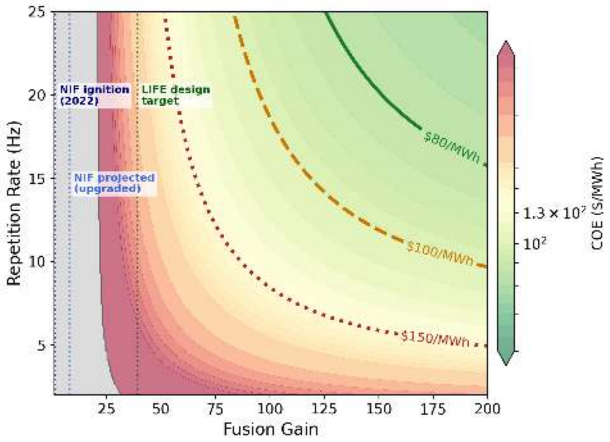

Backup: shadow price — NIF anchor

$\xi = dT^*/d\,\mathrm{COE}_{\mathrm{target}}$: marginal TRL-step^2 per dollar of COE relaxation. Computed by central FD on the outer optimization — a derivative of an *optimal value* through the coupled physics, the same nested-gradient pattern as the robust-design slide.

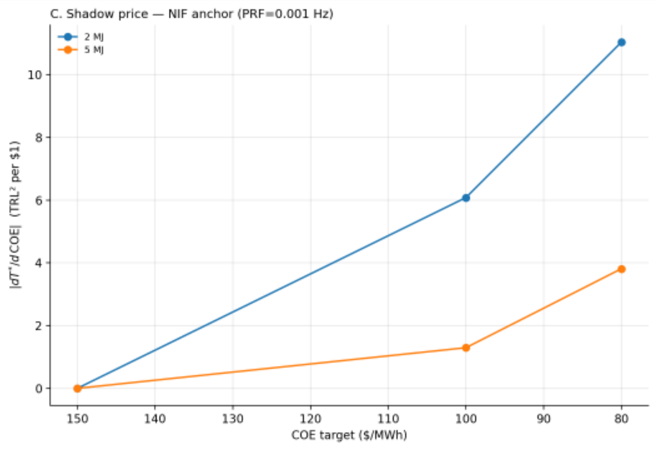

Reading the curve (NIF SOTA = 0.001 Hz):

- LIFE/2 MJ: $|\xi|$ ramps from $\sim$0 at \$150 to $\sim$10 at \$80 — marginal effort multiplies 10$\times$
- HAPL/5 MJ: same shape but 2–3$\times$ flatter — the bigger driver absorbs more of each capability push
- Convexity is the "infeasibility wall": as $\mathrm{COE}_{\mathrm{tgt}}$ approaches the achievable minimum, ROI on extra TRL collapses

Backup: shadow price — Mercury anchor

Same quantity, with rep rate SOTA $= 5\,$**Hz instead of** $0.001\,$**Hz.** "If you grant lab-rep technology, what's the marginal cost of tightening COE?"

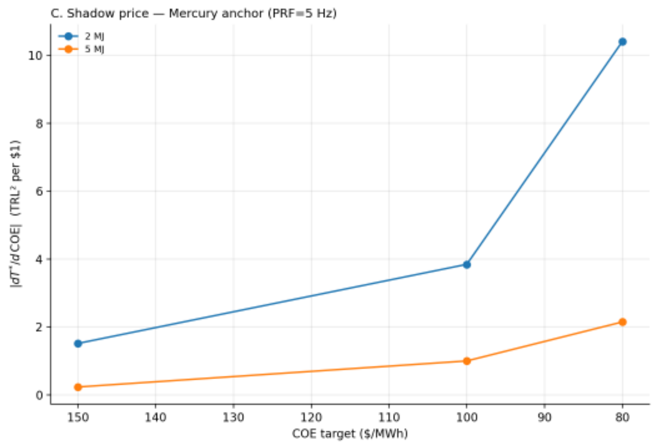

Reading the curve (Mercury SOTA $= 5\,$Hz):

- LIFE/2 MJ peaks at $|\xi| \approx 10$ at \$80 — comparable to NIF anchor at the same target
- HAPL/5 MJ stays under 2.5 across the range
- **Anchor sensitivity at \$150 is large:** NIF/2 MJ shadow price ≈ 0; Mercury/2 MJ already ≈ 1.5. NIF still has free room to spend on rep-rate before COE matters; Mercury has already paid that cost

Backup: Rad-hydro target physics

1D Lagrangian radiation-hydrodynamics solver (SRT):

- Compressible Euler equations (Lagrangian)
- Sector ray-traced laser deposition (64 rays)
- Flux-limited Spitzer thermal conduction
- Electron-ion equilibration
- 128 cells, adaptive time-stepping
- Full JAX — `jax.grad` works end-to-end

Differentiable w.r.t.:

- Laser pulse shape (temporal profile)
- Laser energy
- Flux limiter, Coulomb log

Outputs:

- Fusion gain, yield, burn fraction

Backup: Related work

Work	**Approach**	**Gap we address**
GEM (LLNL)	Excel, 1D sweeps	One-at-a-time sweeps; no simultaneous ∇
pyFECONs (Woodruff 2026)	Python parametric, MC	Sweeps & sampling; no AD
Hawker 2020	14-param IFE, MC	No decomposition; no gradients
Whyte 2026 (Q_{econ})	Analytical criterion	We can compute ∇Q_{econ}
Araiinejad 2025	MCF TEA	MCF only
This work	Differentiable JAX DAG	**Coupled gradients; path decomposition; fidelity-aware; multi-point ranking**

Broader context: Differentiable simulation is increasingly common in weather/climate (CliMA, NeuralGCM), materials, molecular dynamics. We bring the same approach to fusion plant design.

Backup: Key references

- W. Meier et al., "Systems Modeling for the Laser Fusion-Fission Energy (LIFE) Power Plant," *Fusion Sci. Technol.* 56(2), 647–651 (2009)
- GEM: Generalized Economic Model, LLNL-CODE-2013481 (2025) — released IFE techno-economic model
- P. Amendt et al., *Phys. Plasmas* 18 (2011) — empirical gain curve
- R. Miles et al., "LIFE Target Fabrication Costs," LLNL-TR-416932 (2009)
- S. Woodruff et al., "Costing methodology for fusion power plants" (2026) — pyFECONs
- D. Whyte et al., "Economic viability criteria for fusion" (2026) — Q_{econ}
- S. Hawker, "IFE economics model" (2020) — 14-parameter Monte Carlo
- J. Schwartz et al., "Grid value of fusion energy" (2023) — GenX modeling
- MIT Energy Initiative, "The Role of Fusion in Decarbonization" (2024)
- F. Spitzer et al., "Is fusion too late for Europe?" (2025)
- L. Tang et al., "Fusion learning rates are overestimated," *Nature Energy* (2026)
- A. Lindley et al., "Can fusion become cost-competitive?" (2023)
- M. Araiinejad & K. Shirvan, "TEA of magnetic confinement fusion" (2025)