SAJAS Stellarator

Deuteron

By: Soham, Aarya, Jessica, Aayush, Sara

Problem Statement

Optimize and create a business case for a volumetric **neutron source** based on a **quasi-axisymmetric stellarator** design to produce a medical isotope **Molybdenum-99.**

Introduction Mo-99: A crucial element

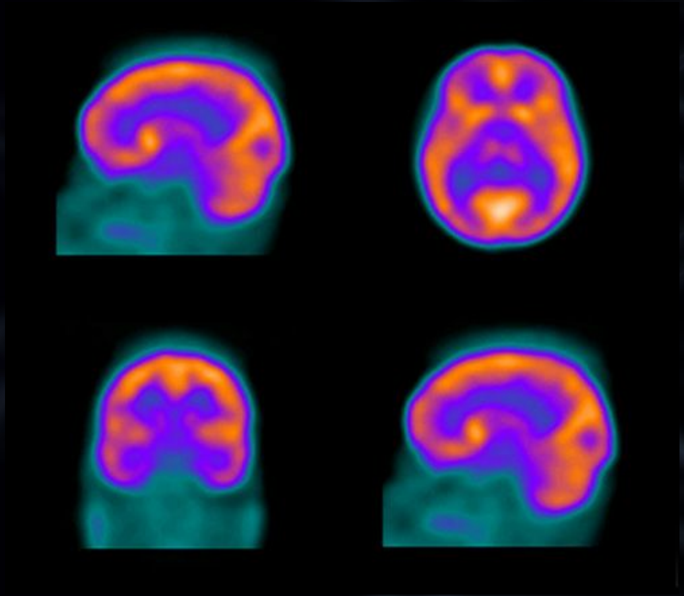

Figure 1: SPECT (single-photon emission computerized tomography scan) on the brain

Medical Applications: Parent isotope of Tc-99m used as radioactive tracer for diagnosing and monitoring medical conditions including **cancer**, **heart conditions**, **brain conditions,** impacting **56,000** Americans every day [1].

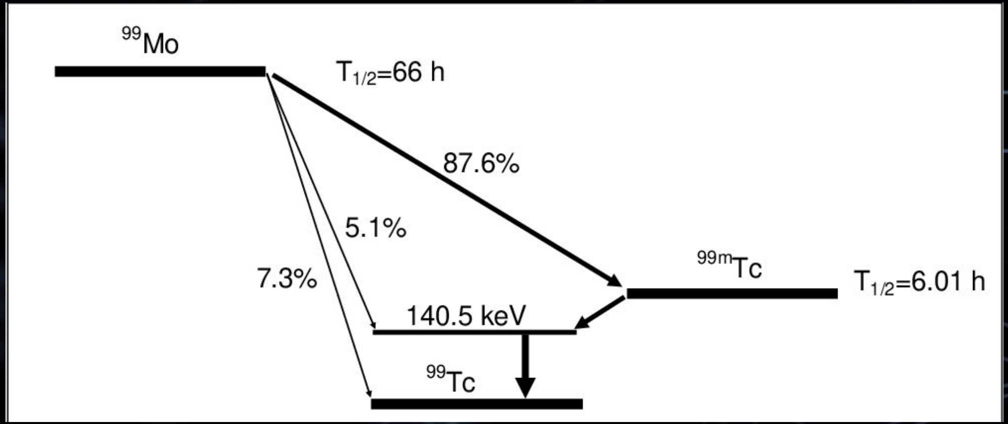

Context: Producing Mo-99

Small government research reactors are the primary producers of Mo-99:

- **NRU, Canada**
- **HFR, Netherlands**
- **BR2, Belgium**
- **SAFARI1, South Africa**

There are 3 ways of producing Mo-99!

1. U-235 to Mo-99

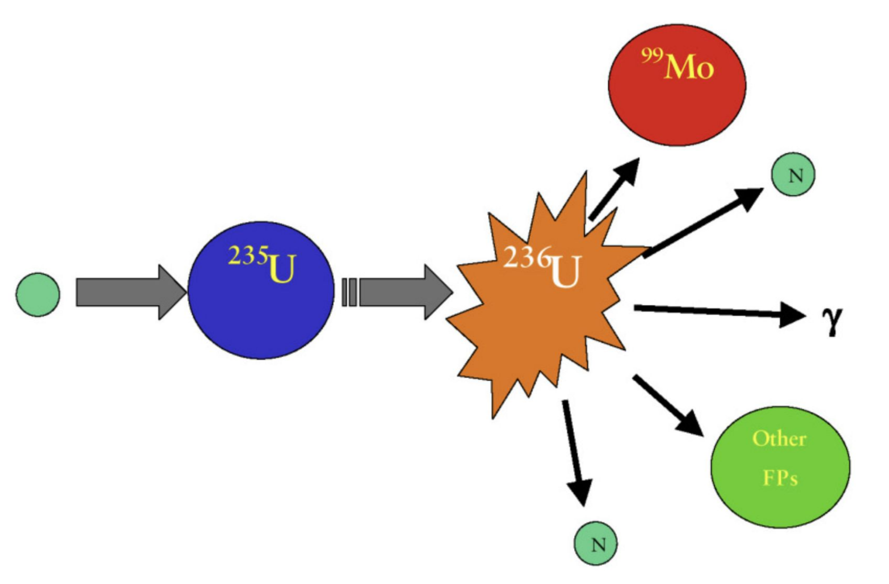

$$^{235}\mathrm{U} + n_{\mathrm{th}} \rightarrow {}^{99}\mathrm{Mo} + {}^{135}\mathrm{Ba} + 2\,n + Q$$

Benefits:

- High thermal cross-section (585 barns)
- High specific activity
- Mo-99 produced is nearly pure, easily separable from U-235

Limitations:

- Reliability of U-235 supply
- Both HEU and LEU are heavily regulated by the IAEA and require safeguard systems
- Large U-235 starting material increases solid and radioactive waste.

2. Mo-100 to Mo-99

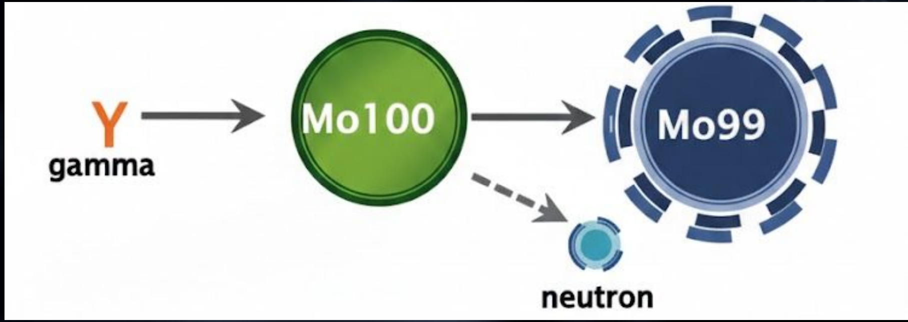

$$^{100}\mathrm{Mo} + n_{\mathrm{fast}} \rightarrow {}^{99}\mathrm{Mo} + 2\,n$$

Benefits:

- No Uranium requirements
- No harmful waste

Limitations:

- Requires D-T fusion to meet threshold of 8 MeV
- Fast neutrons required, increasing price proportional with volume.

3. Mo-98 to Mo-99

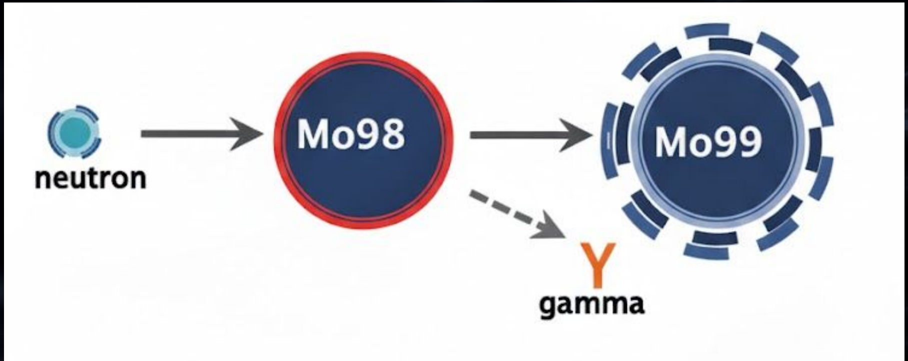

$$^{98}\mathrm{Mo} + n_{\mathrm{th}} \rightarrow {}^{99}\mathrm{Mo} + \gamma$$

Benefits:

- No Uranium requirements
- No harmful waste production
- Compatibility with both D-D and D-T fusion

Limitations:

- Low cross section compared to fission reduces specific activity
- Product is Mo-99 diluted by Mo-98

Mo99 extraction

RadioGenix (automated chromatography)

Already approved by the FDA.

- Instead of static column, use fluidic system
- Advantage: volume of Mo is independent of extraction efficiency. Works effectively regardless of specific activity
- Cost: generator unit is more expensive than disposable lead pot but is reusable

Zirconium Gel Generator

If RadioGenix is too expensive, gel generators are a low cost alternative

- Chemistry: irradiated Mo is chemically converted to Zr molybdate gel
- Matrix: gel acts as a column matrix with high Mo capacity (30% by weight)
- Performance: allows for elution of Tc99m with reasonable concentration even with low specific activity sources (like ours)
- Allows export of product across the world as this elution is low cost and can be done anywhere.

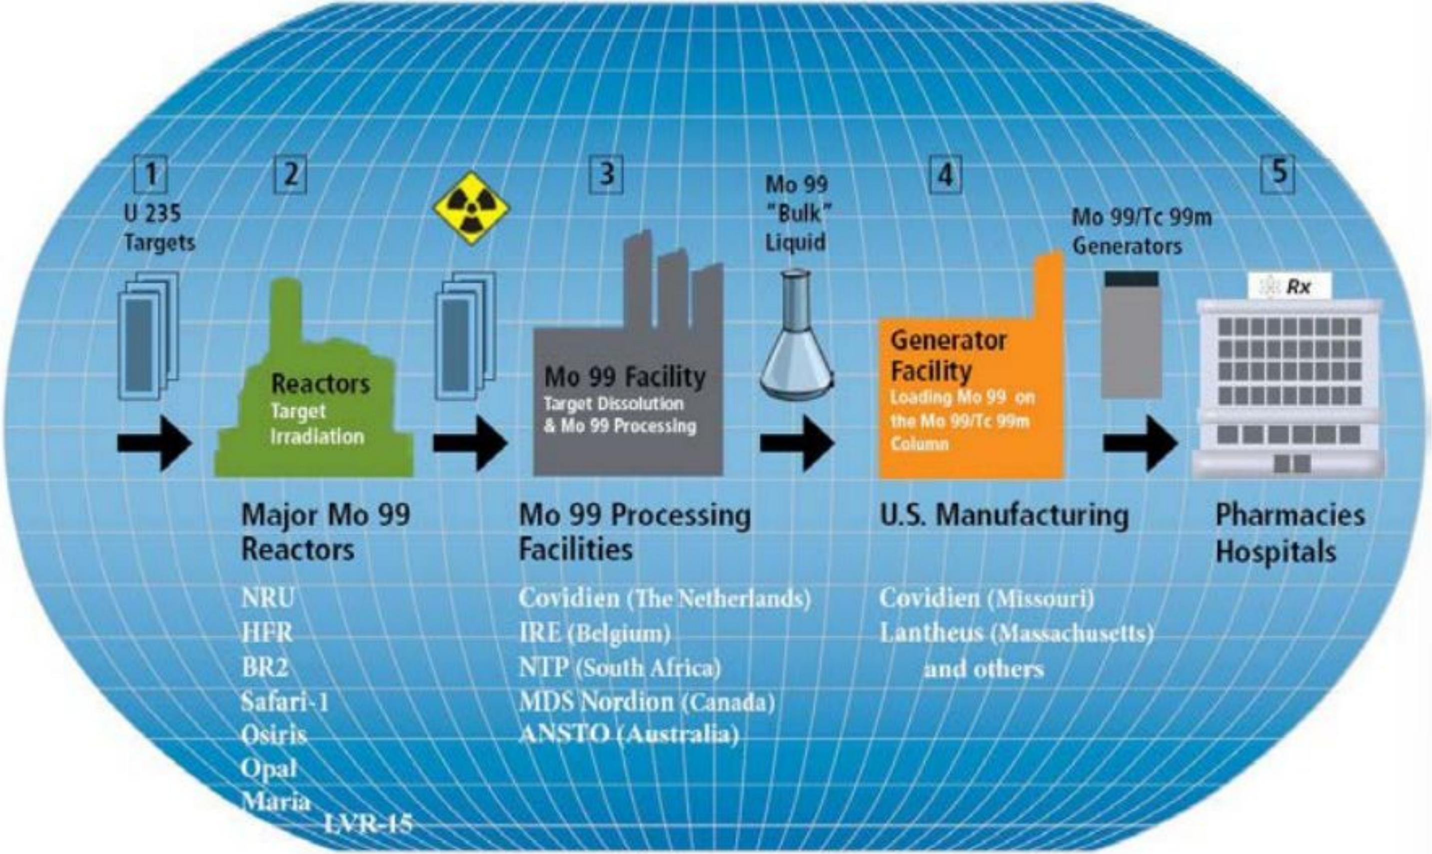

FIG. 3. The global supply chain of ^{99}Mo production and subsequent utilization schematics. Source: www.covidien.com

Stellarator at a Glance

Quasi-axisymmetric (QA): For better confinement

Stat	Value	Meaning
R_0	2.0 m	Major radius
a	0.22m	Minor radius
Aspect ratio	9.1	R/a
ι (iota)	0.10	Rotational transform (ρ = 2/3)
Volume	1.92 m^3	Plasma volume
B_0	0.58 T	On-axis magnetic field
P_ext	10 MW	External heating power
Neutron rate	8.97×10^{13} /s	D–T neutron source

Mo-99:
1100Ci/year
Medical isotope output

Capital cost: $0.71B
FusionHacks cost model

Operating:
$6.0-6.5M/year

Temperature: T ≈ 2.29 keV

"A QA stellarator with $R = 2$ m, $a = 0.22$ m, $\iota \approx 0.1$, passing coil quality ($\max|B\cdot\hat{n}/B| <= 5\times10^{-3}$), D–T neutron rate $\sim 9\times10^{13}\ s^{-1}$ for Mo-99 production, at ~$0.71B capital cost."

- Us

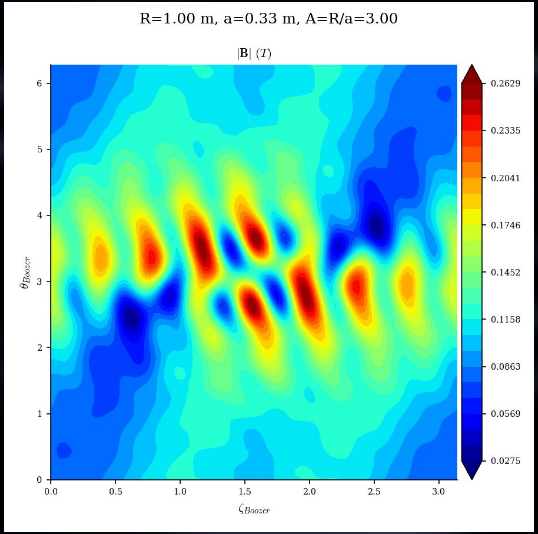

Boozer Coordinates

Boozer coordinates (θ, φ) are a special flux-coordinate system. In quasi-symmetric stellarators, |B| depends mainly on one angle (e.g., φ for QA), not both.

If contours of |B| in (θ, φ) are straight, the field is quasi-symmetric, which reduces neoclassical transport and improves confinement.

https://wiki.fusion.ciemat.es/wiki/Boozer_coordinates#:~:text=Boozer%20coordinates%20are%20a%20set,generality)%20in%20this%20coordinate%20system.

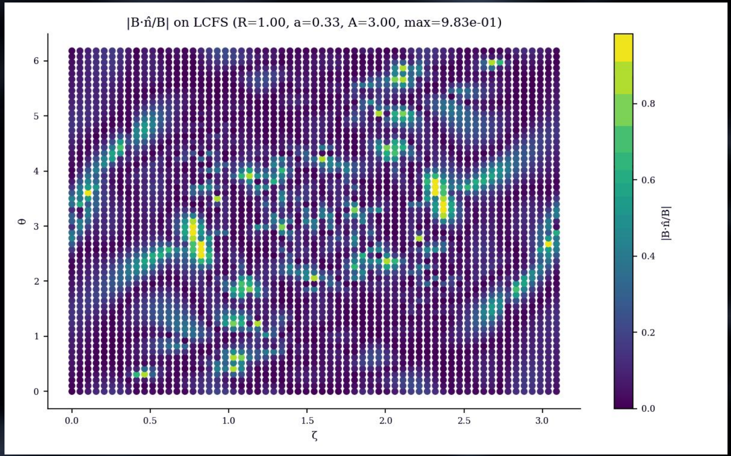

The B field from the coils should stay tangent to the plasma boundary. The normal field error is the component perpendicular to the surface.

$$\frac{\boldsymbol{B}\cdot\hat{\boldsymbol{n}}}{B}\Big|_{\rho=1}$$

- n̂ = outward unit normal on the plasma boundary
- |B·n̂/B| = normalized normal component

Max|B·n̂/B| ≤ 5×10^{-3}

https://www.sciencedirect.com/science/article/pii/S2772828525000159#:~:text=The%20necessary%20magnetic%20field%20for%20plasma%20confinement,target%20surface%20for%20the%20coil%20magnetic%20field.

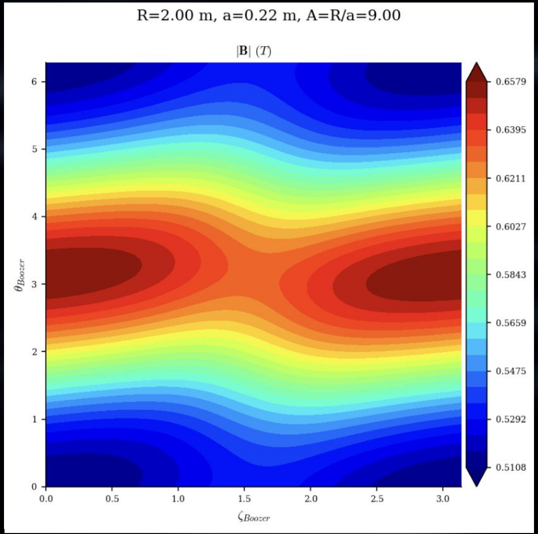

Boozer Coordinates

- Horizontal contours confirm excellent Quasi-Axisymmetry (|B| independent of toroidal angle)

- Limitations include high aspect ratio (A=9) for silliness resulting in a non-compact reactor design.

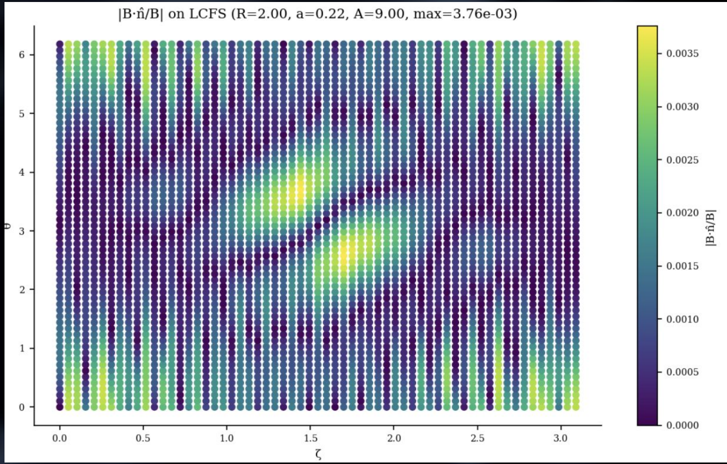

- Normal component of the B field on the LCFS
- This maps surface normal field errors. A max error of ~0.38% is decent, but indicates imperfect flux surfaces causing edge magnetic islands.

$$\frac{\boldsymbol{B}\cdot\hat{\boldsymbol{n}}}{B}\Big|_{\rho=1}$$

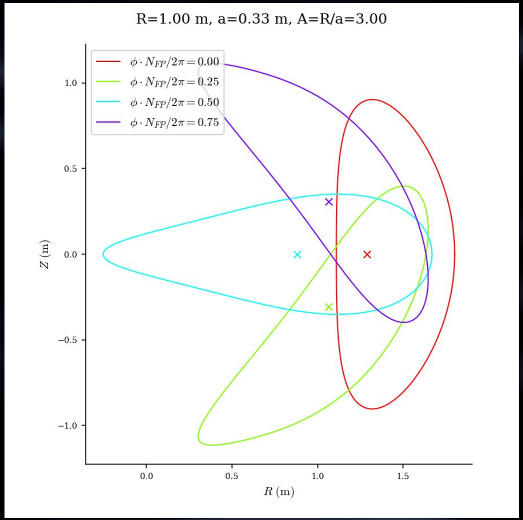

Boundary

- The Last Closed Flux Surface (LCFS) is the outer plasma boundary, defined by the plasma-vacuum interface.
- The shape (elongation, triangularity, etc.) affects confinement, stability, and how easily coils can produce the target field.

https://arxiv.org/pdf/2505.10709#:~:text=The%20LCFS%20represents%20the%20transition%20between%20the,LCFS%2C%20is%20not%20directly%20measured%20during%20experiments.

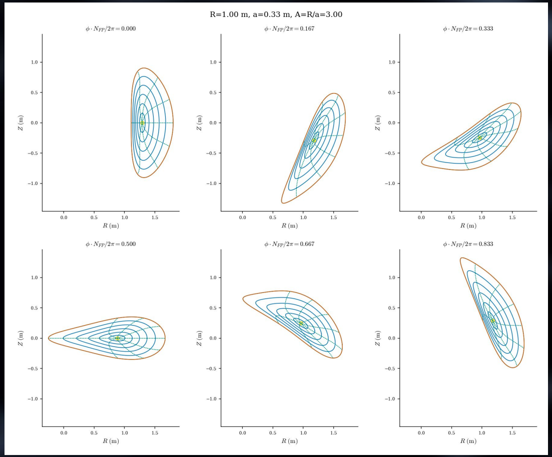

Flux Surfaces

- Magnetic field lines wrap around nested, closed surfaces called flux surfaces.
  - Each surface is a constant value of the flux coordinate (ρ from 0 at the magnetic axis to 1 at the plasma edge).

For good confinement, they should be smooth and closed.

https://en.wikipedia.org/wiki/Flux_surface

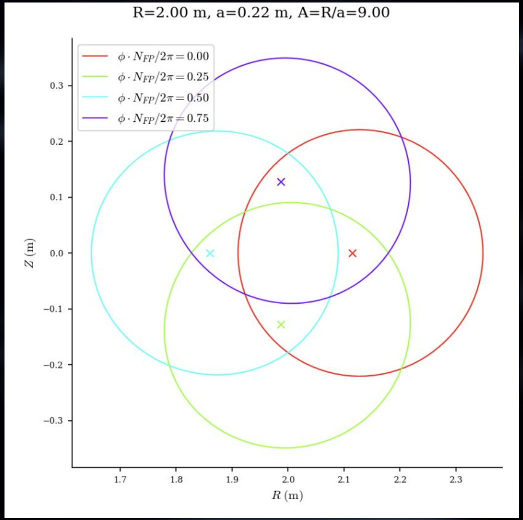

Boundary

- Rotating circular cross-sections
- Design likely suffers from poor confinement, limiting the neutron flux intensity required for efficient Mo-99 production.

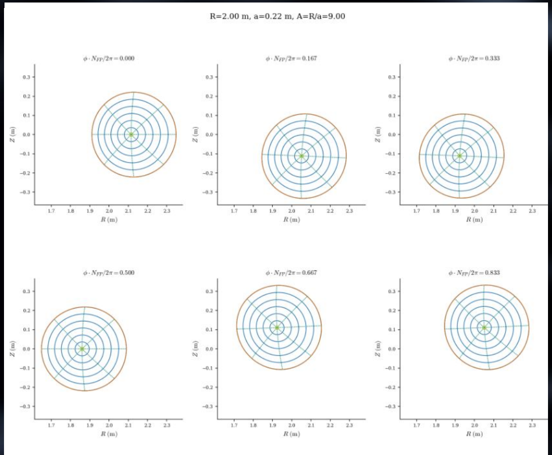

Flux Surfaces

- The plots show nested, circular flux surfaces, confirming good magnetic confinement.
- The lack of shaping restricts plasma pressure, which may limit the neutron flux intensity needed for commercial Mo-99 production rates.

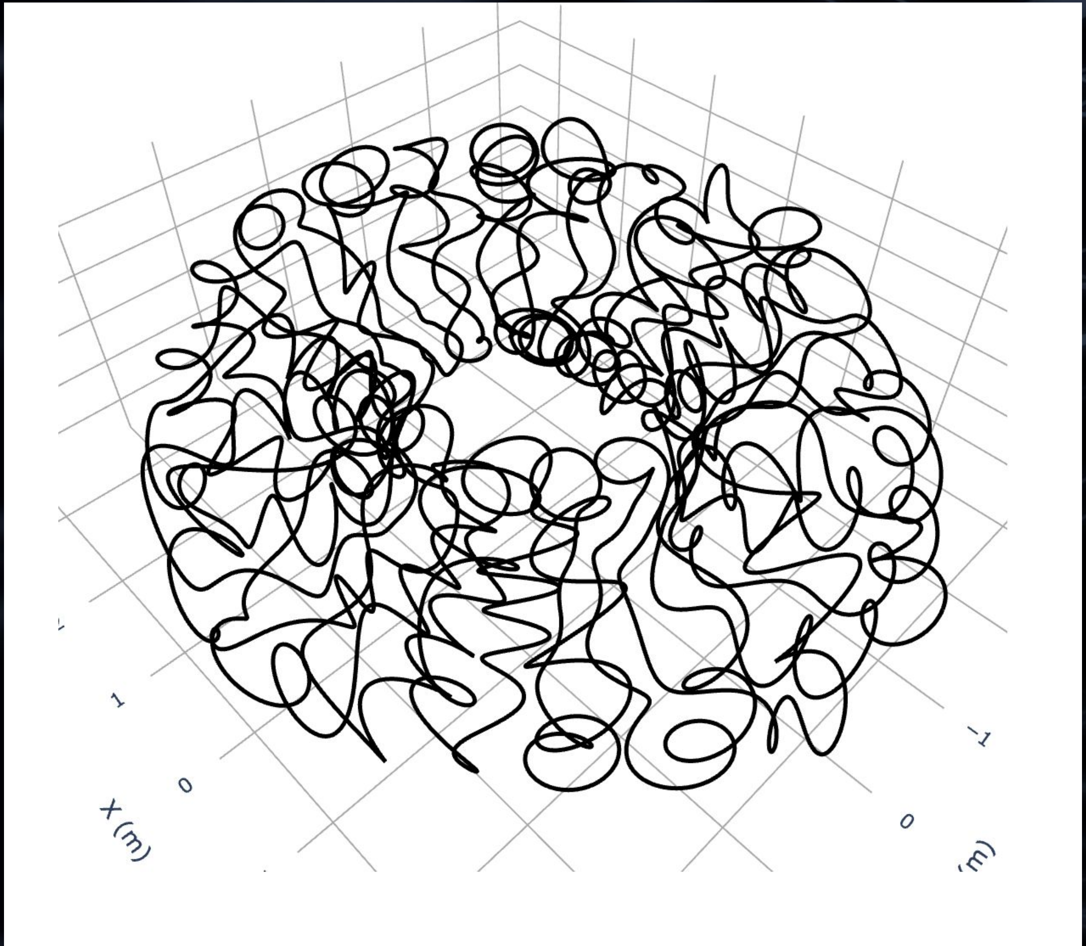

Coils

- Non-planar, twisted modular coils, so harder to manufacture.
- $\iota = 0.1$
- 24 coils
- UGLY but doesn't overlap 💔

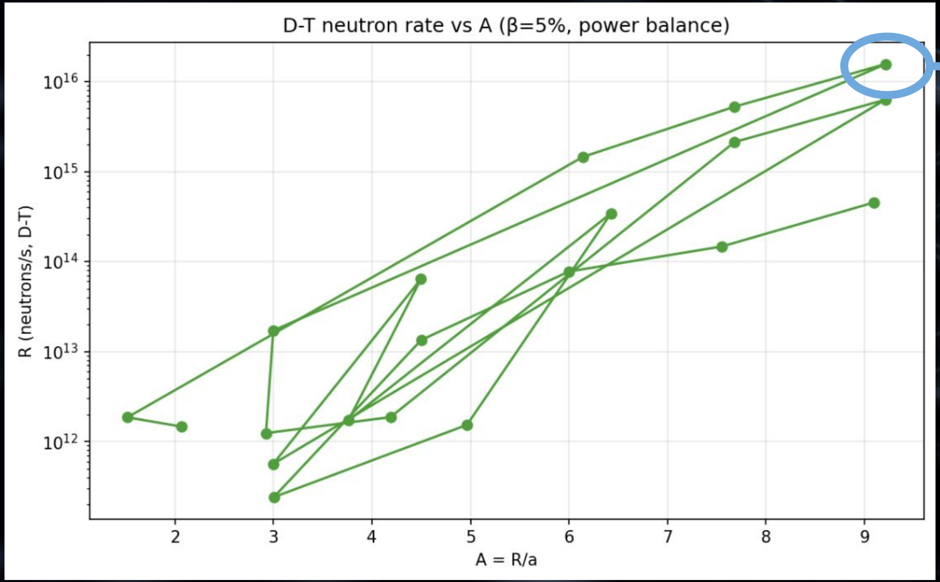
D-T neutron rate vs A (β=5%, power balance)
R (neutrons/s, D-T)
10^{16}
10^{15}
10^{14}
10^{13}
10^{12}
2
3
4
5
6
7
8
9
A = R/a

Us!

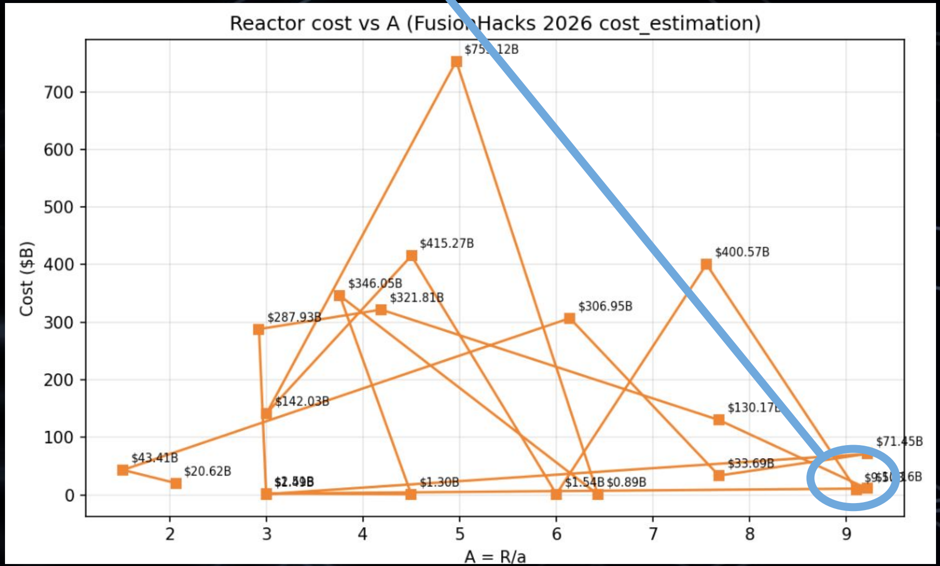
Reactor cost vs A (FusionHacks 2026 cost_estimation)
Cost ($B)
700
600
500
400
300
200
100
0
$43.41B
$20.62B
$287.93B
$142.03B
$346.05B
$321.81B
$415.27B
$1.30B
$306.95B
$1.54B
$0.89B
$400.57B
$33.69B
$71.45B
2
3
4
5
6
7
8
9
A = R/a

Project Parameters

Parameter	Value	Source
Project Life Span	**10 years**	Challenge base case; stellarator components (coils, blanket) may justify longer (20–30 yr) with appropriate maintenance.
Discount rate (WACC)	**10%**	Standard for fusion/cleantech; adjust for risk profile.
Corporate tax rate	**25%**	Challenge specification.
Depreciation	**Per unit over lifespan**	Estimate total Mo-99 capability = Q_Mo99_Ci_year × project life; justify depreciation schedule.
Fixed overhead	**$1M/year**	Challenge specification.
Working capital	**10% of revenue**	Challenge specification; recovered in final year.
Market demand growth	**5%/year**	Challenge specification.

Capacity assumption

Parameter	Value	Source
Capacity factor	**0.8**	FusionHacks: R_a < R for maintenance. IAEA (2024): research reactors optimize availability; 80% allows planned maintenance, refueling, downtime. See References. Sensitivity: 0.7–0.9.

COGS: Electricity (heating)

Parameter	Value	Source
Formula	$P_{ext, MW}$ x 1000 x 8760 x capacity x $/kWh	Annual electricity cost, EIA Table 5.6.A: U.S. industrial 7.85¢/kWh (Nov 2024); LA 5.05¢, RI 22.73¢. See References.
Electricity price	**$0.07/kWh** base	EIA Table 5.6.A: U.S. industrial 7.85¢/kWh (Nov 2024); LA 5.05¢, RI 22.73¢. See References.

COGS: Mo-98 / Mo-100 target

Parameter	Value	Source
Mo-100(n,2n)Mo-99 σ	~1.47 barn @ 14 MeV	IAEA EXFOR/NDS: Kong et al. ~1.5 barn; 1.471 ± 0.31 barn @ 13.59 MeV. See References.
Enriched Mo cost	Quote required	Enriched Mo-98/Mo-100 from commercial suppliers; premium vs natural Mo (USGS ~$55/kg). See References.
η (Ci/neutron)	Geometry-dependent	Yield depends on target mass, flux; challenge: "research conversion efficiency."

Neutron-split parameters

Parameter	Value	Source
f_mo99	0.5	50% Mo-99, 50% tritium; adjust if Q_T < burn rate
TBR	0.25	UofT fusion reactor number

Current Market

- High demand low supply
- **Non-collusive oligopolistic Market (BR-2, HFR and NRU account for >50% market.**
- Current economic structure does not support investments required for new production infrastructure.
- Increase in prices can cause a change of $80 000 - $100 000 on hospitals.

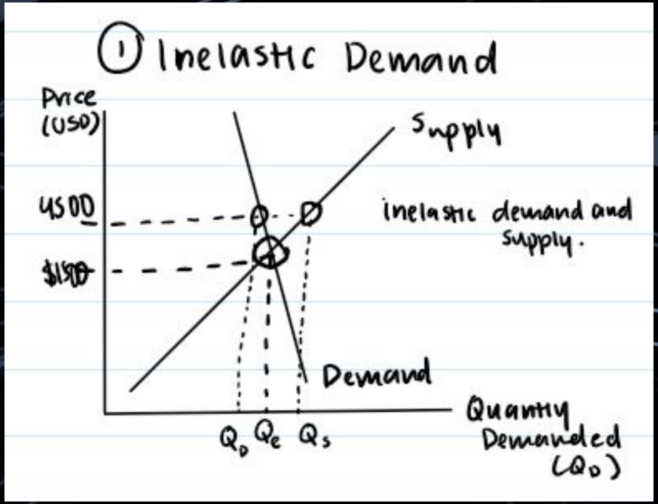

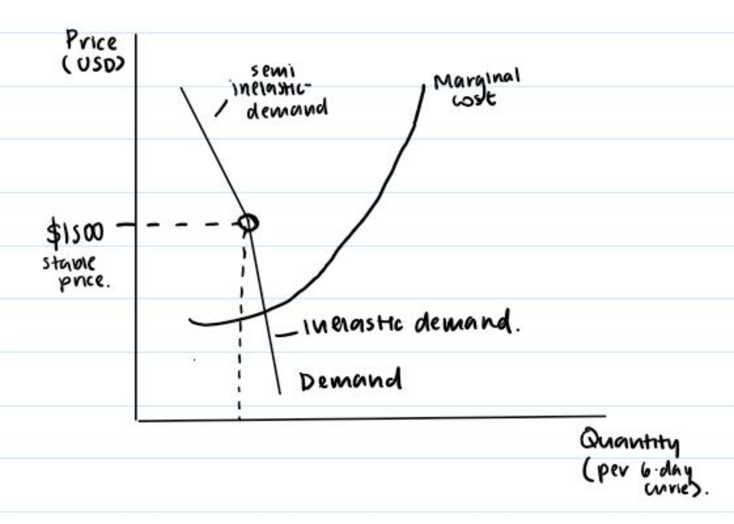

Market Price: $1500 per 6 day curie

Pricing Models

Our Price: $300
(average price = $470)

Marginal Cost

$$(\alpha Q_{\text{Mo-99}}) - 100P = 9000 + P\left(\frac{MC}{1-\frac{1}{N\varepsilon_D}}\right)^{-\varepsilon_D}$$

Supply Elasticity

Number of Competitors

Demand Elasticity

Adding our facility increases competition, can decrease price.

Price of Our Reactor

$$\text{Net Revenue} = (\text{Units Sold} \times \text{Price per unit}) - \text{Returns} - \text{Allowances} - \text{Discounts}$$

The net revenue is deemed to be $9 Million per year. However, as this is an investment, we can create an economy of scale and the price can decrease with time.

Market Failures

- High barriers to entry
  - Massive capital requirements
    - traditional large-scale fusion reactor like ITER: ~22B euros
    - smaller experimental sellarators (Wendelstein 7-X): ~ 370M euros
  - Fission reactor: licensing a fission-based facility takes 5-10 years, costs over $50M
  - Fusion neutron source: D-D fusion licensing time 18-24 months
- Monopoly - limited number of aging nuclear reactors

Market Failures

- Supply Chain Inelasticity
  - Mo-99 has a half-life of ~66 hours
  - Mo-99 declines by 1%/hour from radioactive decay
  - Decay over 6 days reduces Mo-99 to 22% of its initial activity
- **6-Day Curie Problem**

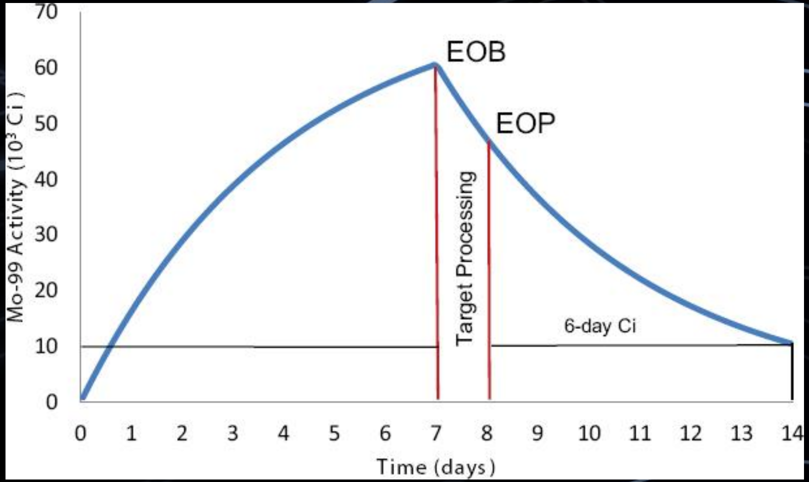

Illustration of the buildup of Mo-99 in a uranium target during irradiation (EOB = end of bombardment), and the decay after processing (EOP = end of target processing)

SWOT Analysis

STRENGTHS

Stellarator requires only external magnetic fields and can run continuously. They also benefit longevity, requiring less fatigue. Medical technology requires constant demand despite prices increases or decreases.

WEAKNESSES

Tokamak is still more proven and yields a higher pivot droppage. Initial stellarator is extremely expensive (10s of billions).

OPPORTUNITIES

Mo-99 is expected to grow from $5.17 - $7.74 billion by 2035, growing 4.6% annually. Aging nuclear reactors will retire, opening space for new competitors. Government subsidies exist for research reactors.

THREATS

Volatility of Mo-98 prices makes production and prices of Mo-99 variable fluctuations in the market.

Competitor Analysis

Company	Price	Production Costs	Quality	Features	Dependencies
SAJAS	$300	$225	Mo-98/Mo-99 mixture output.	No fission → no need for hot cell	Mining and enrichment Mo-98 from BC
MDS Nordion (Partnership with AECL)	$1,500/Ci	~$125-325/Ci (2005)	U-235 leads to high specific activity, Mo-99 easily separable.	LEU-based Single bank of hot cells	Reliant on US & Canadian Governments, U-235 prices
Mallinckrodt	$1400/Ci (2005)			LEU-based 10 hot cells	Largest global supplier (60% global market). Part of larger company (Keenova Therapeutics)
BMS/Lantheus	$2,080/Ci (2005)			LEU-based	Publicly traded company

Projected Capital, Operational and Investment Cost

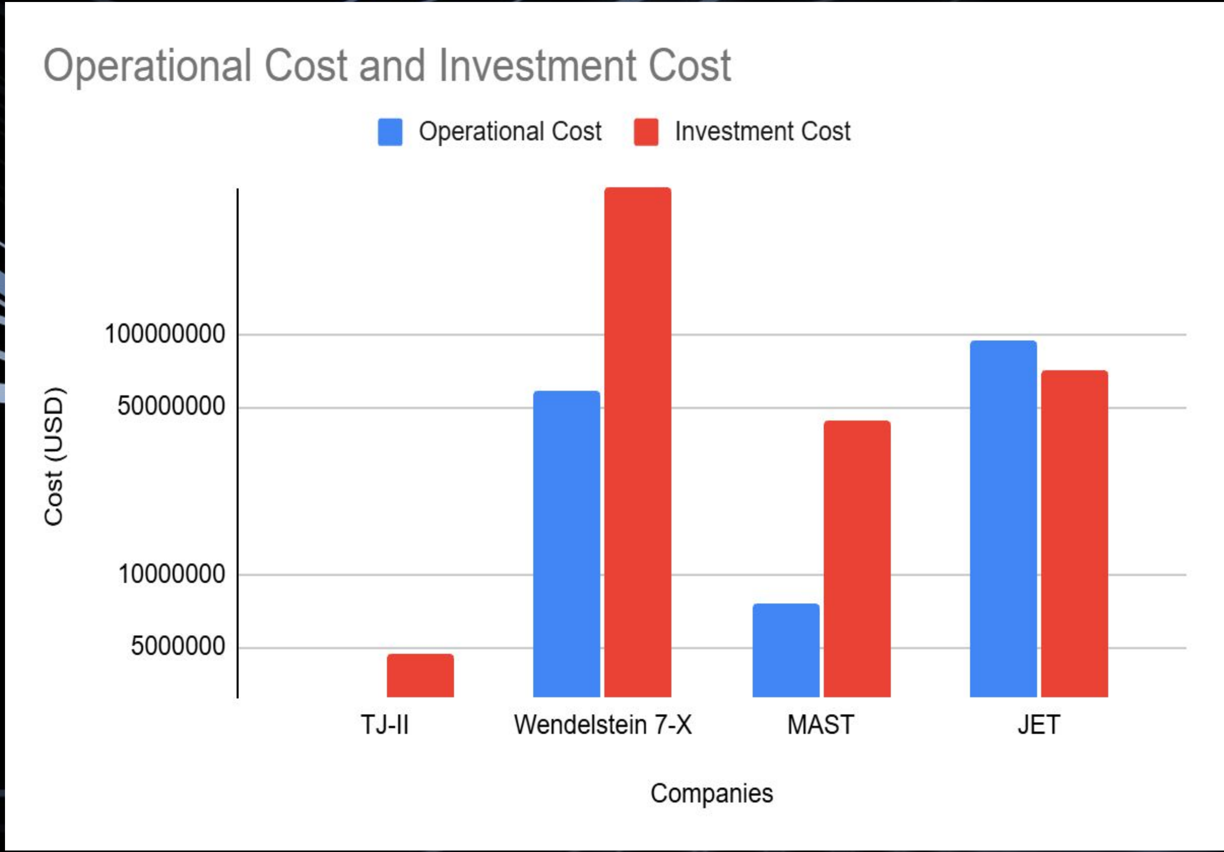

Construction Costs estimated to be **$714 Million**

Operational Costs approx. **$60 Million** ($6 Million for heating and electricity)

Investment Costs approx. **$131 Million**

Depreciation & Taxes

Scientific Research and Experimental Development (SR&ED) tax incentives:

1. Claim a deduction against income
2. Earn an investment tax credit (ITC)

Capital cost acceleration (CCA): Allows manufacturers to write off all the costs of machinery immediately

Depreciation Amount: In 70 years, the depreciation is estimated to be around **$48 Million.**

10+ YEAR TIMELINE

Research → **Investment** → **Operation/Implementation**

~1-2 years

- Consult scientists
- Partner with fusion research facilities
- Develop prototype stellarators to test plasma
- Secure a computing facility (~100 Tflops) capacity

~2-10 years

- Apply for government funds and grants. (NSERC, CIHR, AECL)
- Pitch to investors
- Apply for bank loans
- 2 stages of investing 350 million

~40-50 years

- Testing production facilities
- Large scale maintenance and upgrades every ~10 years
- Adapting to supply chain variations and increases in demand

10+ YEAR TIMELINE: Research

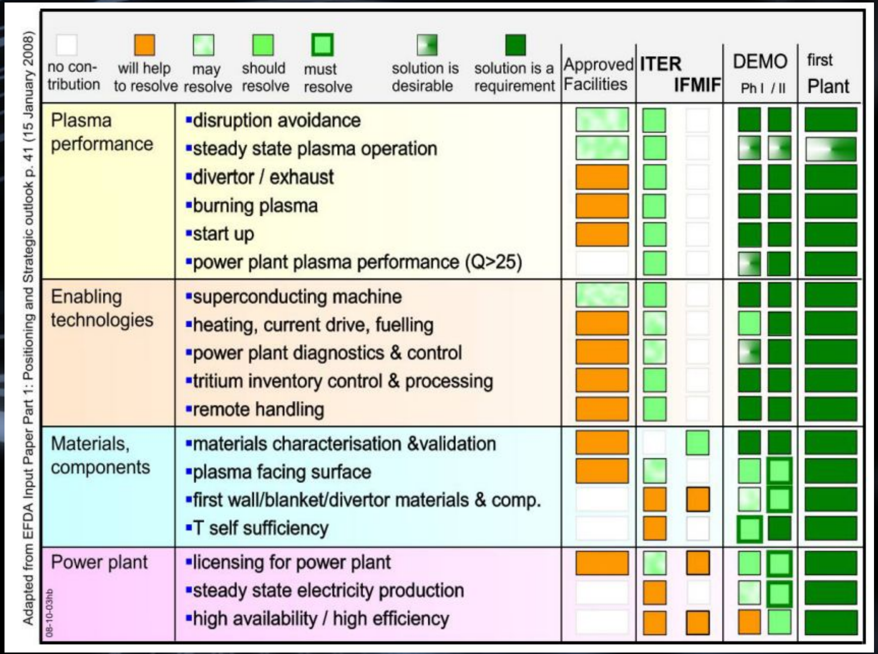

Example: R&D Needs and Requirements for Facilities for Fusion Energy Sources

10+ YEAR TIMELINE: Operations

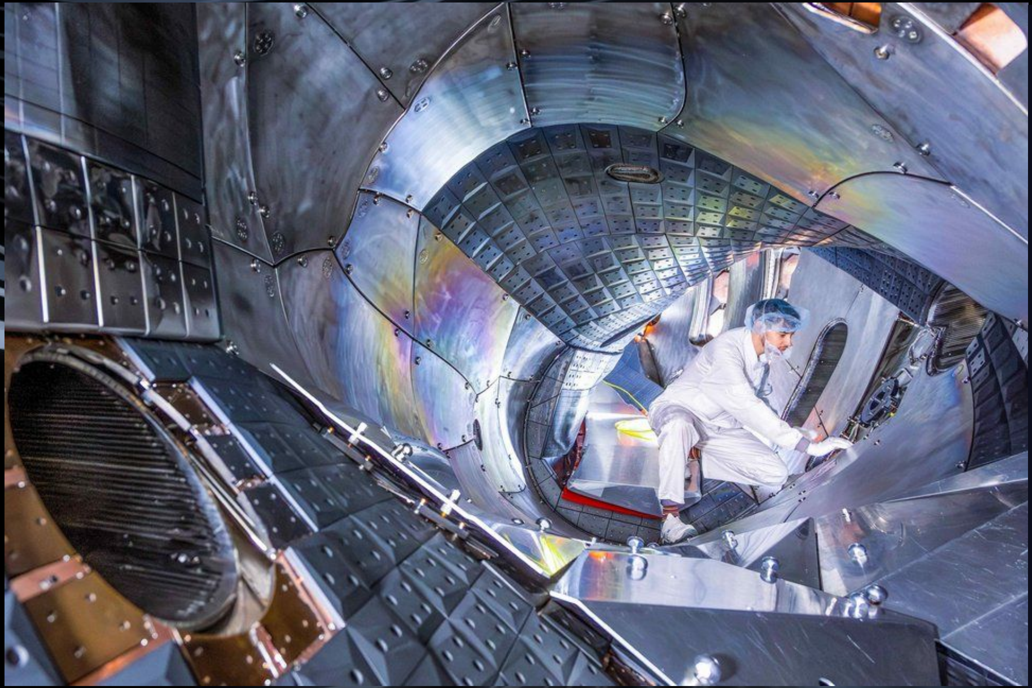

Example: Wendelstein 7-X

- Maintenance every 10 years
- Demand will increase due to increased need for medical tests and devices in aging Canadian population
- Adjust to supply chain changes to ensure Mo-99 reaches target destinations

NPV and Investment Plan

- Estimated starting NPV of -$5-10 billion (based off nuclear reactor references)

- Verdict: Do not recommend this stellarator design as return on investment will not recoup the costs in a reasonable amount of time.

THANK

YOU!

References (research for justification)

Topic	Citation	URL
Electricity (industrial)	U.S. EIA, *Electric Power Monthly*, Table 5.6.A (Industrial, Nov 2024). U.S. total 7.85¢/kWh; LA 5.05¢; RI 22.73¢.	https://www.eia.gov/electricity/monthly/epm_table_grapher.php?t=epmt_5_6_a
Mo-100(n,2n)Mo-99 σ	IAEA EXFOR; Kong Xiangzhong et al., Mo-100(n,2n)Mo-99 @ 13.4–14.8 MeV, ~1.5 barn. Also: 1.471 ± 0.31 barn @ 13.59 MeV (NDS).	https://www-nds.iaea.org/exfor/
Research reactor availability	IAEA, *Optimization of Research Reactor Availability and Reliability: Recommended Practices* (STI/PUB/2080, 2024).	https://www.iaea.org/publications/15504/optimization-of-research-reactor-availability-and-reliability-recommended-practices
Natural Mo price	USGS, *Mineral Commodity Summaries*, Molybdenum (2023 avg $55.60/kg). Enriched Mo requires commercial quote.	https://pubs.usgs.gov/periodicals/mcs2024/mcs2024-molybdenum.pdf
TBR	The Production Rate of Natural Tritium Harmon Craig &Devendra Lal	https://www.tandfonline.com/doi/abs/10.3402/tellusa.v13i1.9430#:~:text=The%20predicted%20production%20rate%20is,from%20an%20extra-terrestrial%20source

Issues With Stellarators

- Manufacturing difficulties and mechanical strength of the coils.
- Irradiation and long-term availability
- Quench protection issues
- Cost of superconductors and cryogenic cooling.
- Breeding blanket breeding the Tritium requires lots of complex engineering and is vulnerable to failure

Confinement time

Implementation: power_balance_solver.py and fusionhacks_metrics.tau_E_iss04(). Power balance uses τ_E to solve for temperature T from:

$$P\_ext + P\_\alpha = W / \tau\_E, \text{ with } \beta = 5\%.$$

Highlight:

- Use of ISS04 stellarator scaling
- Explicit dependencies on R, a, ι, and H
- **Trade-offs:** larger a helps τ_E ∝ a^2.28, but higher A (R/a) can hurt

Stability

Implemented:

- **Mercier** – MercierStability in stage6 (REFINE=1)
  - local MHD stability condition that combines several effects along each flux surface:
- **Ballooning –** BallooningStability in stage6 with STAGE6_BALLOONING=1
  - Instabilities that grow in regions of bad curvature. Perturbations gain energy and can grow.
  - The code uses the infinite-n ideal MHD ballooning equation and solves for the growth rate squared λ

Not implemented:

- **No systematic stability scan or stability constraints in the main optimization loop**

Aspect Ratio

Sweep structure:

- **S**weep of (R, a) with $A = R/a$
- **Baseline:** R = 1.0 m, a from 0.083 m ($A \approx 12$) to 0.52 m ($A \approx 2$)
- **Extra points:** “promising A” (2.0, 2.5, 3.0, 3.5) with R = 1.2 m for fixed A

Optimization:

- **AspectRatio** objective targets desired A
- **RotationalTransform a**t $\rho = 2/3$ targets $\iota \approx 0.42$ ($\tau_E \propto \iota^{0.41}$)
- **Stage 5 plots:** τ_E vs A, cost vs A, $R_{neutrons}$ vs A.

QA Optimization

Definition: Quasi-axisymmetry (QA): B ≈ B(ρ, φ).

Implementation:

- QuasisymmetryTwoTerm(helicity=(1,0)) – primary objective in stage 3, weight 1e2
- Optimized jointly with AspectRatio and RotationalTransform
- QA geometry is fixed after stage 3 and reused in stages 4–6
- In stage 4: only coils are optimized; plasma boundary fixed
- In stage 6 with REFINE=1: QA is kept in the objective while refining ı
- Constraint: The challenge does not recommend optimizing ε (effective ripple); QA is used instead.

Pipeline

Topic	Purpose
Stage 3	R–a sweep; QA optimization (QuasisymmetryTwoTerm); force balance; aspect ratio; ι target → lock geometry
Stage 4	Coil optimization (QuadraticFlux, curvature, length, spacing); target $B\cdot\hat{n}/B \le 5\times10^{-3}$
Stage 5	τ_E, cost, R_neutrons vs A; pick designs; produce geometry_for_planner.json
Stage 6	Select best design; optional ι refinement + Mercier + ballooning; export to final/

Important sequence:

1. QA geometry first, then lock.
2. Coils and physics metrics (τ_E, neutrons, cost) come after.
3. No boundary changes after stage 3.

Future Steps (code)

Stability: Add Mercier/ballooning into the main optimization (stage 3 or 6) instead of only in refinement; run stability checks on a subset of designs.

Confinement: Use the H-factor from ε in all stages rather than H = 1 fallback when possible.

QA metrics: Report effective ripple, QS residuals, or similar as part of the design summary.

Coil simplicity: Add explicit engineering objectives (e.g., curvature limits, maintainability) in stage 4.

shorter runtime!!

Things We Changed

1. ɪ target: 0.42 → 0.1 (stage6, REFINE)

- Pushing ɪ toward 0.42 for better τ_E caused B·n̂/B to exceed 5×10^{-3}. We backed off to ɪ ≈ 0.1 so coils could still pass while improving confinement over very low ɪ designs.

2. Relaxed IOTA_FLOOR to 0

- stage6_finalize_geometry.py line 79 – IOTA_FLOOR = 0.0 with comment "relaxed: pick best B·n̂; prefer higher ɪ when passing".
- Dropping the floor lets us keep passing designs and only prefer higher ɪ when multiple pass.

3. Coil optimization weights: flux dominant, curvature/length relaxed

```
# Flux match paramount; curvature/length heavily relaxed to reach 5e-3
weights = {
    "quadratic flux": 15000,
    "coil-coil min dist": 200,
    "plasma-coil min dist": 50,
    "coil curvature": 50,
    "coil length": 10,
}
```

Things We Changed

4. Coil-coil distance: 8% → 5% of minor radius

- archive/FINDINGS_SUMMARY.md mentions "8% coil-coil; 5% constraint now in code"; current code uses bounds=(0.05, np.inf) in stage4 and stage6.
- Challenge spec requires coil-coil distance ≥ 5% of minor radius. We aligned with that and tightened from the earlier 8%.

5. Aspect-ratio sweep: optional R variation for fixed A

- $\tau_E \propto R^{0.64} a^{2.28}$, so τ_E depends on R and a separately. Varying R at fixed A lets us explore better (R, a) combinations instead of only R=1.

6. RotationalTransform in stage 3 to avoid $\iota \approx 0$

- Pure QA optimization produced $\iota \approx 0.01$–0.02 for high-A designs, killing τ_E ($\propto \iota^{0.41}$). We added a RotationalTransform target at $\iota \approx 0.42$ to keep QA while enforcing useful confinement.

Iota Problems 😭

We tried pushing for higher ı (and lower aspect ratio) to improve confinement (τ_E ∝ ı^0.41), but in most cases the normal field error B·n̂/B went above the 5x10^{-3} limit.

Design type	A (R/a)	$ı_{2/3}$	τ_E (ISS04)	max B·n̂/B	Result
High-A, small a	8–12	0.01–0.02	~10^{-4} s	0.002–0.003	Coils pass
Mid-A	4–6	≈0.42	0.002–0.005 s	0.011–0.026	Exceed limit
Low-A, large a	~2–3	≈0.42	0.01–0.02 s (best τ_E)	0.22–0.99	FAR over limit

QA optimization at low A drove ı toward ~0.42, giving good τ_E.

- Those shapes were harder for coils to reproduce; the normal component of B at the boundary got large.

High-A shapes (small, tight plasmas) gave lower B·n̂/B and passed, but had very low ı because of the QA solution, which hurt confinement.

We avoided a design at the best-τ_E point (A ≈ 2–3) because of coil failure. Instead we selected a **high-A** design (R=2 m, a=0.22 m, A≈9) with passing coils and low ı, refined ı moderately to 0.1 (instead of 0.42), re-optimized coils with QA and ı fixed.

“If it works it works”

Iota Problems 😭

What to do differently next time

- Treat ι, A, and B·n̂/B together
  - Add B·n̂/B (or coil feasibility) as a constraint in the QA / aspect-ratio sweep, not only in the coil step. Avoid designs that are known to fail coil matching.
- Joint plasma–coil optimization
  - Allow boundary and coils to co-evolve instead of locking the plasma and optimizing coils only. Could reveal shapes that give higher ι and acceptable B·n̂/B.
- Explore ι vs coil complexity
  - Cary ι over a range (e.g. 0.05–0.3) and track B·n̂/B for each to map the ι–B·n̂/B trade-off before committing to a final design.
- Multi-objective optimization
  - Pareto fronts over (τ_E, ι, B·n̂/B, cost) to choose a design that balances confinement and coil feasibility rather than optimizing ι alone.