Eight ways to make fusion easier — and which two are worth a week

Cross-disciplinary experimental hypotheses generated from a 120-paper corpus, each cited to primary literature and each novelty claim tested by a recorded adversarial search rather than asserted.

CORPUS	SCREENED FROM	LEVERS COVERED	HYPOTHESES
120 papers 80 core · 40 donor	**2,967** candidates 6 sources, 90 queries	**23/23** 2 thin (n=2)	**8** 6 novel · 1 partial · 1 blocked

What the corpus says before any hypothesis

Three findings constrain everything below. Two are discouraging, and they bear directly on the most popular parts of the design space.

The non-equilibrium p-^{11}B route has a serious negative result
`dist_000` reworks Rider's 1997 analysis with updated cross sections: the recirculating power needed to hold electrons cooler than ions exceeds the *fusion power*, for both Maxwellian and non-Maxwellian electrons. `dist_005` sets the optimistic bound from the other side. Any hypothesis in that gap must name **which of Rider's assumptions it breaks** — and the donor set adds a cost Rider omitted: self-generated waves relax a non-thermal distribution even without collisions.

The avalanche is real but small
Careful treatments put alpha-driven chain multiplication at **1–10%** (`reac_007`, `ener_008`), not the large factor claimed in the contested laser literature. `ener_003` finds the widely reported laser p-^{11}B alpha yields come from beam fusion that does *not* scale to net gain. Treat the avalanche as a bonus, never as a load-bearing assumption.

Polarized fuel is in better shape than its reputation
The two historical objections — polarization won't survive, and you can't make enough of it — have each been retired since they were raised, in *different literatures that do not cite each other*. This is the most underrated lever in the set.

The hypotheses

Ordered as ranked, not as numbered. Every card states what would kill it.

H3 Point fusion's own depressed-collector expertise at alpha energy conversion

lever 12 · alpha extraction | donor: microwave vacuum electronics | computational → benchtop | 12–24 mo | low → moderate cost

Applying multistage depressed collector (MDC) design methodology to fusion alpha direct conversion should push efficiency from the ~80% now demonstrated toward the ~90% routine in traveling wave tubes — and solve the thermal-balance problem fusion DEC has not addressed.

Two communities converged on the same number without citing each other. `ener_012` reports ENN's electrostatic direct conversion at **>80%**, framed as addressing "the current shortage of first-hand experimental data." `denc_006` reports a genetic-algorithm-optimized collector at **88.5% recovery with 0.72% reflux** — designed, fabricated and measured. The problems are structurally identical: a chaotic beam with wide velocity spread arrives at a electrode stack and must be sorted by energy without back-streaming.

Bridge changed this hypothesis. `"depressed collector" AND "fusion"` returns **127** papers — which looked like a kill until I read them. All 127 are gyrotrons: MDCs for ECRH heating systems, at MW CW power, with E-vB variants and 3D secondary-electron modelling. `"depressed collector" AND "alpha particle"` returns **0**. So the capability already exists inside fusion labs and has never been pointed at the other side of the machine. The ask is organizational, not cross-disciplinary.

Kills it: required electrode voltages for MeV alphas exceed practical standoff, or space charge at reactor alpha currents defeats energy sorting. `ener_012` reports direct evidence against the second.

H2 Test lattice screening with NMR — no nuclear measurement at all

levers 3, 18 · screening | donor: solid-state NMR | benchtop | 6–12 mo | low cost | key assumption high-risk

If the conduction-electron density at the deuteron site in ZrD_2 is high enough to produce a 200–300 eV screening potential, that density must also show up as an anomaly in the deuteron spin-lattice relaxation rate — measurable on a benchtop spectrometer.

The corpus contains a flat contradiction: `barr_009` measures $U_e \approx 205{\pm}35$ eV in ZrO_2; `barr_018` reports **<30 eV** on the same material. Both are peer-reviewed. Both infer screening by extrapolating a bare-nucleus S-factor. Meanwhile `cond_004` (1985) shows conduction-electron screening must be included to fit deuteron relaxation — an entirely independent experimental channel onto the same quantity, which the screening literature never invokes.

Bridge: `"spin-lattice relaxation" AND "electron screening" AND "deuteride"` → **1** hit, and it is `cond_004` itself. Nothing in forty years has joined these.

Kills it: a calculation showing the density sampled by hyperfine/quadrupolar coupling can differ from the density relevant at the nuclear surface by the order of magnitude needed to reconcile 200 eV with ordinary relaxation rates. *That calculation may well kill it — but it is cheap, and it it survives, the experiment is weeks on equipment most universities own.*

H5 Polarized fuel's two objections are both answered — in literatures that don't cite each other

levers 9, 12, 14, 16 | donor: polarized-target physics, molecular beams | tokamak campaign | 24–36 mo

Reassembling four separate results makes polarized D-^{3}He in *magnetically* confined geometry the best-supported item in the preexisting stack.

Survival: `reac_013` finds wave-driven depolarization "surprisingly weak" in SPARC/ITER-scale devices, contrary to prior expectation — with one specific caveat, that lower-hybrid heating destroys polarization above $\delta B_0/B_0 = 10^{-3}$ (`reac_004`). That is a distribution, not a refutation. **Production:** `disp_008` runs a polarized target by triplet-DNP at **0.3 T and 100 K** instead of dilution-refrigerator conditions; `reac_008` projects >90% polarization at $>10^{21}$ s^{-1}. **Geometry:** `reac_012` shows hydrodynamic instabilities depolarize implosions — so magnetic geometry is favoured and inertial is not, a distinction the corpus never draws. **Payoff:** `reac_009` gives >3× fusion power; `reac_003` adds that polarization biases alpha perpendicular, turning them into a channeling-wave source.

Bridge found partial prior art: `"dynamic nuclear polarization" AND "fusion fuel"` → **1** hit, "Deuteron polarization of solid DT." Someone has polarized solid DT by DNP. That is a precedent to build on; the conjunction with survival, geometry and the alpha-channeling synergy is still unmade.

Kills it: achievable polarized throughput falling more than three orders of magnitude short of fuel consumption rate. *Caveat:* `reac_009`'s >3× is explicitly optimistic by its authors' own statement — treat as an upper bound.

H1 Settle the screening dispute with the Trojan Horse Method

levers 3, 18, 1 | donor: nuclear astrophysics | accelerator campaign | 18–30 mo

Measuring d+d in a deuterated metal by Trojan Horse Method alongside a conventional direct measurement isolates the screening potential without *extrapolation* — resolving whether U_e is ~200 eV or ~30 eV.

Both sides of the contradiction enter screening by extrapolating a bare-nucleus S-factor, so the approach-dependent extrapolation model. `dast_006` describes the technique built to remove exactly that dependence: quasi-free breakup delivers the participant at ultra-low energy with no entrance-channel Coulomb barrier, giving the **bare-nucleus** S-factor directly. `dast_001` states THM explicitly overcomes "extrapolation procedures and enhancement effects due to electron screening."

Bridge: `"Trojan horse method" AND "electron screening"` → 64, all astrophysical. `AND "metal"` → **0**. `AND "deuterated"` → **0**. The condensed-matter transfer is unclaimed.

Kills it: THM and direct measurement agreeing within combined uncertainty — in which case there is no anomalous lattice screening and levers 3 and 18 lose their central empirical motivation. Also killed if quasi-free breakup kinematics can't be cleanly separated in a dense lattice.

H6 Two lever-11 schemes are underpriced because they ignore wave-driven relaxation

levers 11, 14 | donor: space plasma physics | computational only | 3–6 mo | very low cost | negative result

Analyses of cool-electron/hot-ion operation account for *collisional* thermalization but not for relaxation driven by self-generated waves. Including it makes the non-equilibrium ledger worse than `dist_000` already reports.

`dist_019` claims a p-^{11}B operating point at T_e=600 keV with T_e/T_i=65 keV. Rider in our Rider models what space physics measures routinely: a distribution with enhanced tails destabilizes waves that relax it *without collisions*. `dnes_005` gives whistler heat-flux thresholds; `dnes_007` derives a distinct "halo instability" that grows specifically when there is a halo on the ion tail. An empirical caution sits alongside: `dist_000` found a strongly non-Maxwellian JT-60U plasma gave no reactivity gain.

Bridge: `"halo instability" AND "anenutronic"` → **0**. The 4 hits for whistler+fusion are laser-plasma heat conduction, not reactor power ledgers.

Kills it: collisional relaxation dominating wave-driven relaxation by orders of magnitude at reactor densities, making the addition irrelevant. *Worth writing down even though discouraging — a negative that saves effort on two rungs of the stack is worth as much as a positive.*

H8 Revisit muon catalysis as a lattice problem, not a plasma problem

levers 4, 18, 5, 20 | donor: exotic-atom physics | muon beam facility | 36–48 mo | muon economics likely decisive

The condensed-matter environment is an active variable in muon catalysis that the modern lattice-confinement literature ignores entirely.

`dexo_009` (1993) shows that in a solid D-T admixture, resonant muonic molecule formation acquires a **broad spectral component spanning the solid's elementary-excitation bandwidth** — the lattice absorbs the excess energy, permitting quasi-resonant formation far outside the isolated resonance. `barr_001` independently reports an "anomalous condensed-matter effect." Meanwhile every lattice-confinement paper in the corpus is about electron screening and never touches muons, which screen by a factor of a magnitude harder. `barr_017` states the gap: 150 fusions/muon achieved, **~20x more needed**, with alpha sticking the binding loss — and `barr_012` identifies the density-dependent stopping power, a *material* property, as what controls it.

Bridge: `"muon catalyzed fusion" AND "lattice confinement"` → **0**. AND `"host metal"` → **0**.

Kills it: sticking varying by less than 20% across accessible host materials and temperatures. *Note the standing objection this does not address:* muon production cost is probably the real binding constraint, and this hypothesis improves only the physics ratio.

H4 Recirculate a cooled, resonance-tuned proton beam through a dense boron target

levers 15, 2, 1, 11 | donor: accelerator physics | desk calc → ring time | 6 mo calc | probably fails

A storage ring running a phase-space-cooled proton beam at energy tuned to the reactant boron target breaks Rider's objection differently: the non-thermal distribution is maintained by *external cooling against scattering*, not by recirculating power against collisional thermalization.

`dacc_005` describes HESR at FAIR — cooled beams through **dense internal pellet targets**, momentum spread to 10^{-5}, beam equilibria worked out. That is the architecture. And `dacc_004` (IOTA) supplies the number that decides it: beam lifetime under 7 minutes, fewer than 10^9 turns, losses itemized. So the question is sharp arithmetic — **does a proton make enough fusion energy in ~10^9 passes to pay for its cooling and replacement?**

Bridge: `"storage ring" AND "proton boron"` → **0**. `"beam cooling" AND "anenutronic"` → **0**. The 3 hits for electron-cooling-fusion-reactor are divertor plasma cooling — a different sense of the word.

Kills it: required cooling power exceeding fusion power at any achievable target areal density — which would simply reproduce Rider's conclusion in different coordinates. **Stop rule:** if the first arithmetic pass shows cooling power exceeds fusion power by more than 10×, do not proceed to ring time.

H7 Pre-excite the target nucleus in a storage ring — blocked

levers 8, 15, 23 | storage ring + electron target | blocked on contested foundation

If nuclear excitation by electron capture is confirmed, a storage ring is the natural place to exploit it — the same machine holds a cooled beam *and* provides the electron target.

`reac_019` computes NEEC isomer triggering in storage rings and finds it more efficient than bremsstrahlung or NEET. But NEEC may not exist. `denc_008` (Nature 2018) reported first observation at P_{null} = 0.01; `denc_003` cannot reproduce it under any scenario; `denc_008` (2025) calls confirmation "elusive."

Bridge added a third refutation not in the corpus: "Possible overestimation of isomer depletion due to contamination." Three independent challenges to one Nature claim.

Do this first, and it isn't fusion: run `denc_009`'s decision gate — ^{235}U in laser-cluster interaction engineered so NEEC accounts for >99.9% of yield, removing the attribution ambiguity that sank the ^{93m}Mo claim. Until that exists, this is not fundable.

Novelty, tested

Every claim is the result of a recorded query on OpenAlex, not a judgement; arXiv is not used: it returns 0 on bridges where OpenAlex finds dozens of prior-art papers, so an arXiv zero carries no information.

ID	DECISIVE BRIDGE QUERY	HITS	VERDICT
H1	`"Trojan horse method" AND "metal"`	0	novel
H2	`"spin-lattice relaxation" AND "electron screening" AND "deuteride"`	1	novel
H3	`"depressed collector" AND "alpha particle"`	0	novel, reshaped
H4	`"storage ring" AND "proton boron"`	0	novel
H5	`"dynamic nuclear polarization" AND "fusion fuel"`	1	partial prior art
H6	`"halo instability" AND "anenutronic"`	0	novel
H7	`"nuclear excitation by electron capture" AND "fusion"`	3	blocked
H8	`"muon catalyzed fusion" AND "lattice confinement"`	0	novel

The one that changed my mind was **H3**: a 127-hit bridge looked like a kill and turned out to strengthen the hypothesis, because reading the titles showed the prior art sits on the heating side of the machine. A count alone would have been actively misleading — which is the argument for reading bridge results rather than thresholding them.

What is deliberately not here

Vacuum-modification schemes (lever 22)
The donor papers are genuinely good — `dcm_000` demonstrates the dynamical Casimir effect in a Josephson metamaterial; `dcm_005` reaches deep-strong photon-phonon coupling. But the energy scales are 10–12 orders of magnitude from nuclear, no hypothesis here bridges that gap, and I am not going to manufacture one. The four lowest-quality papers in the entire corpus are in this family.

Antimatter ignition (lever 21)
`ener_015` and `ener_014` are competent, but `dacc_009` documents what an industrial antiproton complex actually costs to run, and nothing makes production economics tractable.

Phonon-driven barrier modulation (lever 19)
The Floquet donor evidence is strong — `dqct_002` and `dqct_005` show tunneling rates experimentally controlled, including full suppression, with a specific spacetime-symmetry requirement. The lattice side is where the corpus's worst papers live. A serious version needs to show the drive frequency can approach the relevant energy scale, which is exactly what those preprints assume rather than demonstrate. **Highest-risk, highest-upside gap in the brief.**

If you had…

BUDGET	GO THIS	BECAUSE
One week, desk only	H6, and the H4 arithmetic	Both are pure calculation and both are disconfirming — they tell you where not to go.
A year and a benchtop	H2, then H1 if H2 survives	Cheapest decisive test of the screening anomaly that motivates two whole levers.
One funded programme	H3	The only hypothesis where both halves are already demonstrated experimentally and the work is engineering transfer, not new physics.

Method. Three query sets against OpenAlex, arXiv, INSPIRE-HEP, OSTI, Semantic Scholar and NASA ADS. Core queries find problems; donor queries find mechanisms and deliberately exclude fusion terms; a donor query mentioning fusion returns only connections already made. Bridge queries run last, to kill or sharpen rather than to inspire. Every paper carries an `evidence_class` — *established / contested / single-group / theoretical* — assigned on the evidence rather than on the topic's reputation; the corpus deliberately retains both a Nature paper claiming the first observation of NEEC and the calculation that cannot reproduce it. Full pipeline, queries and scoring rubric in the repository.