AIR Journal of Engineering and Technology, Vol. 2026
Article Unique ID: AIRJET2026736
DOI: 10.65737/AIRJET2026736

Non-Equilibrium Quantum Thermodynamics of Nuclear Fusion: Resource-Theoretic Bounds, Fluctuation Theorems, and Thermodynamic Uncertainty Relations for Plasma Ignition

Mosab Hawarey
Director, Geospatial Research (*director@geospatial.ch*)

Received: 2026 May 10 | Revised: 2026 May 13 | Approved: 2026 May 14 | Published: 2026 May 15

Abstract

Fusion plasma physics and quantum thermodynamics have developed over several decades as mature but essentially disconnected disciplines, with virtually no cross-citations linking the two literatures. We bridge this gap by establishing the first unified theoretical framework — Non-Equilibrium Quantum Thermodynamics of Fusion (NEQT-Fusion) — and prove four theorems that impose new, fundamental constraints on fusion plasma ignition and power balance.
Theorem 1 (Resource-Theoretic Lawson Criterion). We reformulate ignition as a state-conversion problem in the resource theory of thermodynamics. A family of generalized free-energy conditions, $F_\alpha(\rho) \geq F_\alpha(\rho_{\text{ign}})$ for all Rényi orders $\alpha \geq 0$, is shown to be necessary for ignition. These conditions reduce to the classical Lawson triple-product criterion in the macroscopic Maxwellian limit but provide strictly stronger constraints for non-Maxwellian plasmas. We construct an explicit proton–boron-11 (p-^{11}B) counterexample in which the classical Lawson criterion is satisfied yet a higher-order Rényi condition ($\alpha = \infty$) is violated, rendering ignition thermodynamically forbidden.
Theorem 2 (Quantum Crooks Relation for Fusion Reactivity). The ratio of forward (fusion) to reverse (disassembly) stochastic work distributions for plasma-screened Coulomb tunneling satisfies the exact quantum Crooks fluctuation theorem, $P_F(W)/P_R(-W) = \exp[\beta(W - \Delta F)]$. This yields a thermodynamic ceiling on non-Maxwellian reactivity enhancement via the Jarzynski equality and motivates new, model-independent diagnostics: the Crooks asymmetry curvature κ_R as a non-equilibrium indicator and the Jarzynski-extracted athermality from neutron time-of-flight spectra in NIF-class burning plasmas.
Theorem 3 (TUR Bound on Recirculating Power). The precision of maintaining a non-Maxwellian ion distribution against collisional relaxation is bounded by total entropy production through the thermodynamic uncertainty relation, $\mathrm{Var}(P_{\text{rec}})/\langle P_{\text{rec}}\rangle^2 \geq 2k_B/\dot{S}_{\text{tot}}$. In the quantum-degenerate electron regime, this bound is sharpened, reinforcing the energetic impossibility of steady-state aneutronic fusion via a precision–dissipation tradeoff absent from classical analyses.
Theorem 4 (Unified NEQT-Fusion Criterion). The preceding three results are combined into a single master inequality governing the net engineering gain Q_{eng}, subject to the simultaneous satisfaction of all generalized free-energy conditions, the reactivity ceiling, and the TUR power-balance constraint.
The framework subsumes the Lawson criterion (1957) and Rider's recirculating-power bound

(1997) as limiting cases, while furnishing strictly new constraints in non-Maxwellian, mesoscopic, and quantum-degenerate regimes. It thereby establishes NEQT-Fusion as a new interdisciplinary research direction at the intersection of quantum information science and fusion energy science.

Keywords: quantum thermodynamics; nuclear fusion; resource theory; fluctuation theorems; thermodynamic uncertainty relations; Lawson criterion; non-Maxwellian plasma; Rényi divergence; thermo-majorization; recirculating power

PACS / Subject Codes: 52.55.-s (Magnetic confinement and equilibrium); 52.57.-z (Inertial confinement fusion); 03.65.-w (Quantum mechanics); 05.70.Ln (Nonequilibrium thermodynamics)

Citation: Hawarey, M. (2026). Non-Equilibrium Quantum Thermodynamics of Nuclear Fusion: Resource-Theoretic Bounds, Fluctuation Theorems, and Thermodynamic Uncertainty Relations for Plasma Ignition. *AIR Journal of Engineering and Technology*, Vol. 2026, AIRJET2026736. https://doi.org/10.65737/AIRJET2026736

1. Introduction

1.1 The Quest for Controlled Fusion and the Lawson Criterion

The pursuit of controlled thermonuclear fusion as a terrestrial energy source has been guided for seven decades by a single, remarkably durable inequality. In 1957, Lawson derived the minimum density–confinement-time product $n\tau_E$ required for a deuterium–tritium (DT) plasma to achieve net energy gain, establishing the triple-product figure of merit $nT\tau_E$ that remains the principal yardstick for reactor progress [1]. Extensions by numerous authors have refined the criterion to account for impurity radiation, profile effects, and different confinement geometries, and the recent achievement of ignition at the National Ignition Facility (NIF) in December 2022 has been assessed precisely in terms of how the experimental triple product crosses the Lawson threshold [2, 3].

Yet the Lawson criterion and all its standard extensions rest on a set of assumptions that are known to be violated — sometimes severely — in real fusion plasmas. Chief among these is the assumption of local thermodynamic equilibrium: that all ion species are described by Maxwellian velocity distributions at a single, well-defined temperature T_i. The present work demonstrates that relaxing this assumption, using the modern tools of quantum thermodynamics, reveals an entire family of additional necessary conditions for ignition that the classical criterion cannot detect.

1.2 The Non-Equilibrium Problem in Fusion Plasmas

Evidence that fusion plasmas deviate from Maxwellian equilibrium is abundant and growing. Hartouni et al. [4] recently reported direct spectroscopic evidence for supra-thermal ion populations in NIF burning plasmas, confirming theoretical expectations. Molvig et al. [5] showed that preferential consumption of fast ions in the Knudsen layer of an inertially confined hot spot can deplete the high-energy tail of the distribution by up to 50%, substantially reducing the thermonuclear burn rate. Alpha-channeling schemes, pioneered by Fisch and Rax [6], deliberately exploit non-Maxwellian alpha-particle distributions to channel fusion-born energy into fuel-ion heating. Beam-driven and radio-frequency-heated plasmas in magnetic confinement devices routinely sustain non-Maxwellian populations for performance optimization.

The consequences of non-equilibrium for fusion feasibility were analyzed in a landmark pair of papers by Rider [7, 8], who demonstrated that any fusion system operating with a non-Maxwellian ion distribution must continuously supply recirculating power to maintain that distribution against collisional relaxation. For the aneutronic fuel cycle p-^{11}B, Rider showed that this recirculating power vastly exceeds the fusion power output, making steady-state net energy production impossible under his assumptions. This conclusion has been revisited recently by Guo et al. [9] and Binderbauer et al. [10], who confirmed Rider's central result while identifying narrow parameter windows at elevated electron temperatures. The Rider bound, however, is derived entirely within the framework of classical kinetic theory and uses only mean-value power balances. It does not exploit the full structure of modern thermodynamics.

1.3 The Quantum Thermodynamics Revolution

Over the past two decades, the field of quantum thermodynamics has undergone a transformation driven by the resource-theory approach to thermal processes. Brandão, Horodecki, Ng, Oppenheim, and Wehner [11] established that the second law of thermodynamics is not a single inequality but a *family* of inequalities — the "many second laws" — parametrized by the Rényi order $\alpha \geq 0$. Each generalized free energy $F_\alpha(\rho)$ must decrease under thermal operations, and only when *all* conditions are simultaneously satisfied is a state transformation thermodynamically permitted. This framework subsumes the classical second law (the $\alpha = 1$ case) while providing strictly additional constraints for nanoscale, mesoscopic, and non-equilibrium systems [12, 13, 14].

In parallel, quantum fluctuation theorems have extended the classical results of Jarzynski [15] and Crooks [16] to driven quantum systems. The quantum Crooks relation, rigorously established by Campisi, Hänggi, and Talkner [17] and given a streamlined derivation by Cohen and Imry [18], relates the ratio of forward and reverse work distributions to the free-energy difference via an exact exponential identity. Åberg [19] extended these results to fully quantum settings with coherence.

A third pillar — thermodynamic uncertainty relations (TURs) — was introduced by Barato and Seifert [20], who proved that the relative fluctuation of any steady-state current is bounded below by the reciprocal of the total entropy production. This fundamental precision–dissipation tradeoff has been extended to driven systems by Koyuk and Seifert [21], to open quantum systems by Hasegawa [22], and to systems with multiple conserved quantities by Dechant [23].

These three classes of results — resource theories, fluctuation theorems, and TURs — have transformed our understanding of thermodynamics at small scales and far from equilibrium. They have found applications in biophysics, nanoscale engines, and quantum computing. They have *never* been applied to fusion plasmas.

1.4 The Gap: An Empty Citation Bridge

We emphasize that the cross-citation graph between quantum thermodynamics and fusion plasma physics is, at the time of writing, essentially empty. A systematic search reveals no published work that applies the resource theory of thermodynamics, the many second laws, or thermodynamic uncertainty relations to fusion reactor physics.

The closest existing bridge is the program of Diaz-Torres and collaborators at the University of Surrey, who developed the coupled-channels density matrix (CCDM) formalism for low-energy heavy-ion fusion [24, 25]. Lee and Diaz-Torres demonstrated that quantum coherence between nuclear excited states can enhance sub-barrier fusion probabilities by 15–37%, a genuinely quantum-thermodynamic effect. However, this work addresses individual nuclear reactions at astrophysical energies, not reactor-scale plasmas or power-balance constraints.

There also exists a growing DOE-funded program on Quantum Information Science for Fusion Energy Science (QIS-FES) [26], but its focus is on quantum *computing* algorithms for plasma simulation — quantum hardware for classical problems — rather than on identifying quantum *physics* that constrains fusion thermodynamics. The present work fills the complementary role: quantum thermodynamic *theory* applied to fusion *physics*.

1.5 Our Contribution

We establish the Non-Equilibrium Quantum Thermodynamics of Fusion (NEQT-Fusion) framework by proving four theorems:

Theorem 1 (Section 3) recasts ignition as a state-conversion problem in the resource theory of thermodynamics. We show that the transition from a pre-ignition state ρ_{cold} to a burning state ρ_{ign} requires $F_\alpha(\rho_{\text{cold}}) + W \geq F_\alpha(\rho_{\text{ign}})$ for every Rényi order $\alpha \geq 0$ — a continuum of conditions that collapse to the single classical Lawson inequality only in the macroscopic Maxwellian limit (Theorem 1a). We construct an explicit p-^{11}B counterexample demonstrating that the $\alpha = 1$ (classical) condition can be satisfied while the $\alpha = \infty$ condition is violated, establishing a wider impossibility window for aneutronic fusion than Rider's analysis can detect (Theorem 1b).

Theorem 2 (Section 4) proves the quantum Crooks fluctuation theorem for plasma-screened Coulomb tunneling. The exact identity $P_F(W)/P_R(-W) = \exp[\beta(W - \Delta F)]$ holds for the bimodal work distribution of a single fusion event (reflection peak + fusion peak), yielding a thermodynamic ceiling on reactivity enhancement (Proposition (Reactivity Ceiling)) and two new model-independent diagnostics for burning plasmas: the Crooks asymmetry curvature (Proposition (Crooks Asymmetry Diagnostic)) and the Jarzynski-extracted athermality (Proposition (Jarzynski Athermality Extraction)).

Theorem 3 (Section 5) applies the thermodynamic uncertainty relation to the recirculating power current, showing that the precision of maintaining a non-Maxwellian distribution is bounded by total entropy production: $\mathrm{Var}(P_{\text{rec}})/\langle P_{\text{rec}}\rangle^2 \geq 2k_B/\dot{S}_{\text{tot}}$. Quantum corrections in the degenerate-electron regime sharpen this bound (Theorem 3, quantum-sharpened form), reinforcing Rider's impossibility conclusion through an independent, fluctuation-theoretic argument.

Theorem 4 (Section 6) combines all three results into a unified criterion: four simultaneous necessary conditions — resource-theoretic ignition, reactivity ceiling, precision–dissipation tradeoff, and net power — must be satisfied for a non-equilibrium fusion reactor to produce net energy. The resulting master inequality subsumes the Lawson criterion (1957) and Rider's bound (1997) as limiting cases while providing strictly new constraints in every regime where the plasma departs from macroscopic thermal equilibrium.

The framework thus achieves three things. First, it *subsumes* the foundational results of classical fusion theory, proving backward compatibility. Second, it *extends* them with rigorously derived, physically motivated additional conditions that are non-trivial for non-Maxwellian, mesoscopic, and quantum-degenerate plasmas. Third, it *proposes* new experimental diagnostics — the Crooks asymmetry test and the Jarzynski athermality extraction — that are in principle applicable to NIF-class burning plasmas with existing or near-future neutron spectroscopy capabilities.

1.6 Paper Organization

The remainder of this paper is organized as follows. Section 2 establishes the quantum-thermodynamic foundations: the system–bath partition for a fusion plasma, the Hamiltonian structure including screening, the generalized free energies and Rényi divergence spectrum, the catalog of non-Maxwellian distributions and their athermality, and the decoherence regime analysis that guarantees universal validity of the framework across all fusion regimes. Section 3 develops the resource-theoretic ignition criterion (Theorem 1), including the thermo-majorization formulation, the many second laws applied to plasma state conversion, the classical Lawson recovery, and the constructive p-^{11}B counterexample. Section 4 derives the quantum Crooks relation for fusion reactivity (Theorem 2), including the stochastic work distribution for plasma-screened tunneling, applications to DT and p-^{11}B, and the fluctuation-theorem diagnostics for burning plasmas. Section 5 establishes the TUR bound on recirculating power (Theorem 3), recasts Rider's framework in stochastic thermodynamics, derives the quantum-sharpened bound, and analyzes implications for aneutronic feasibility. Section 6 presents the unified NEQT-Fusion criterion (Theorem 4), the logical hierarchy of constraints, all limiting cases, and the compact master inequality. Section 7 discusses the relation to existing work, limitations, and future directions. Section 8 concludes.

2. Preliminaries — Quantum Thermodynamic Foundations

This section establishes the theoretical apparatus required for the four main theorems. We define the system–bath partition for a fusion plasma (§2.1), characterize non-equilibrium distributions using resource-theoretic monotones (§2.2), and analyze the decoherence landscape to identify where quantum corrections are most significant (§2.3). The presentation is designed to be self-contained for readers from either the quantum thermodynamics or the fusion plasma physics community.

2.1 Fusion Plasma as a Quantum Thermodynamic System

2.1.1 System–Bath Partition

We model a fusion plasma as an open quantum system by partitioning it into a **system** S — the reacting ion pair undergoing tunneling and nuclear interaction — and an **environment** E — the surrounding plasma comprising all other ions, electrons, and collective electromagnetic modes. The system Hilbert space factorizes as

$$\mathcal{H}_S = \mathcal{H}_{\text{rel}} \otimes \mathcal{H}_{\text{nuc}}$$

where $\mathcal{H}_{\text{rel}}$ is the space of relative-motion wave functions (continuum states plus compound-nucleus resonances) and $\mathcal{H}_{\text{nuc}}$ captures internal nuclear-state channels (entrance channel, compound nucleus, exit channels). The environment Hilbert space decomposes as

$$\mathcal{H}_E = \mathcal{H}_{\text{ions}} \otimes \mathcal{H}_{\text{electrons}} \otimes \mathcal{H}_{\text{modes}}$$

encompassing background ions (density n_i, temperature T_i), electrons (density n_e, temperature T_e), and collective modes (plasmons, ion-acoustic waves). The composite state ρ_{SE} lives on $\mathcal{H}_S \otimes \mathcal{H}_E$, and the system state is obtained by partial trace: $\rho_S = \text{Tr}_E[\rho_{SE}]$.

2.1.2 Hamiltonian Structure and Screening Regimes

The total Hamiltonian decomposes as

$$H = H_S + H_E + H_{\text{int}}$$

The system Hamiltonian for the relative motion of two nuclei with reduced mass μ and charges $Z_1 e$, $Z_2 e$ is

$$H_S = \frac{p^2}{2\mu} + V_C(r) + V_N(r)$$

where $V_C(r) = Z_1 Z_2 e^2/(4\pi\epsilon_0 r)$ is the bare Coulomb potential and $V_N(r)$ is the short-range nuclear potential. The Coulomb barrier height $V_B = Z_1 Z_2 e^2/(4\pi\epsilon_0 R_B)$ evaluates to approximately 0.4 MeV for DT (at $R_B \approx 3.6$ fm) and 2.1 MeV for p-^{11}B (at $R_B \approx 3.5$ fm), reaching ≈ 2.7 MeV at closer approach ($R_B \approx 2.7$ fm).

The system–environment coupling enters through the screened potential $H_{\text{int}} = V_{\text{screen}}(r; n_e, T_e) + V_{\text{coll}}$, which replaces the bare Coulomb barrier with an effective potential $V_{\text{eff}}(r) = V_C(r) + V_{\text{screen}}(r)$. Three screening regimes are relevant to the present framework:

1. **Weak (Debye) screening.** $V_{\text{screen}} \approx -V_C(r)\left[1 - e^{-r/\lambda_D}\right]$, where $\lambda_D = \sqrt{\epsilon_0 k_{\text{B}} T_e/(n_e e^2)}$ is the Debye length. This applies in the low-density, high-temperature limit characteristic of magnetic confinement (MCF) plasmas.

2. **Ion-sphere (strong-coupling) screening.** $V_{\text{screen}} \approx -(3Z_1Z_2e^2)/(8\pi\epsilon_0 R_{\text{WS}})$, where $R_{\text{WS}} = \left(3/(4\pi n_i)\right)^{1/3}$ is the Wigner–Seitz radius. This applies in dense, strongly coupled plasmas ($\Gamma \gtrsim 1$).
3. **Quantum (Lim–Son) screening.** A full density-matrix treatment of the electron response yields screening enhancements up to an order of magnitude beyond the classical Salpeter formula in intermediate coupling regimes [27]. This is the regime where the quantum-thermodynamic treatment of the bath is essential.

The residual collisional coupling V_{coll} represents ion–ion Coulomb collisions that drive the system toward thermal equilibrium at rate ν_C.

2.1.3 Thermodynamic Potentials and Generalized Free Energies

At inverse temperature $\beta = 1/(k_{\text{B}}T)$, the Gibbs (thermal equilibrium) state of the system is

$$\gamma_\beta = \frac{e^{-\beta H_S}}{Z(\beta)}, \qquad Z(\beta) = \text{Tr}\left[e^{-\beta H_S}\right]$$

For an arbitrary system state ρ_S, the non-equilibrium free energy is

$$F(\rho_S) = \langle H_S\rangle_{\rho_S} - k_{\text{B}}T\,S(\rho_S) = k_{\text{B}}T\,D\left(\rho_S \parallel \gamma_\beta\right) + F_{\text{eq}}$$

where $S(\rho_S) = -\text{Tr}[\rho_S \ln\rho_S]$ is the von Neumann entropy, $D(\rho \parallel \sigma) = \text{Tr}[\rho(\ln\rho - \ln\sigma)]$ is the quantum relative entropy, and $F_{\text{eq}} = -k_{\text{B}}T\ln Z(\beta)$. The quantity

$$\mathcal{A}(\rho_S) \equiv D\left(\rho_S \parallel \gamma_\beta\right) \geq 0$$

is the **athermality** — the thermodynamic distance of the state from equilibrium — and equals zero if and only if $\rho_S = \gamma_\beta$.

The resource theory of thermodynamics [11, 12] introduces a one-parameter family of **generalized free energies**:

$$F_\alpha(\rho_S) = k_{\text{B}}T\,D_\alpha\left(\rho_S \parallel \gamma_\beta\right) + F_{\text{eq}} \qquad \alpha \geq 0$$

where D_α is the quantum Rényi divergence:

$$D_\alpha(\rho \parallel \sigma) = \frac{1}{\alpha - 1}\ln\text{Tr}[\rho^\alpha \sigma^{1-\alpha}]$$

The central result of Brandão et al. [11] is that a state transformation $\rho_S \to \rho_S{}'$ is achievable under thermal operations only if $F_\alpha(\rho_S) \geq F_\alpha(\rho_S{}')$ for **all** $\alpha \geq 0$ simultaneously. This family of "many second laws" subsumes the classical second law ($\alpha = 1$) and provides strictly additional constraints whenever the system is not in the macroscopic, near-equilibrium regime.

2.1.4 Dictionary: Quantum Thermodynamics ↔ Fusion Plasma

Table 1 provides a complete mapping between the quantum-thermodynamic formalism and standard fusion plasma quantities.

Table 1. Correspondence between quantum thermodynamic and fusion plasma quantities.

Quantum thermodynamic quantity	Fusion plasma equivalent
System state ρ_S	Velocity distribution of reacting ion pair
Gibbs state $\gamma\beta$	Maxwellian distribution at temperature Ti
Athermality $\mathcal{A}(\rho_S) = D(\rho_S \parallel \gamma_\beta)$	Degree of non-Maxwellianity
Non-equilibrium free energy $F(\rho_S)$	Available work for maintaining non-thermal distribution
Generalized free energy $F_\alpha(\rho_S)$	α-order thermodynamic potential (no classical analogue)
Inverse temperature β	$1/(k_\mathrm{B}T_i)$
System energy $\langle H_S \rangle$	Center-of-mass kinetic energy of colliding pair
Partition function $Z(\beta)$	Normalization of Maxwellian at T_i
Bath coupling constant	Coulomb collision frequency ν_C
Confinement time τ_E	Inverse rate of energy loss to external sink

The fusion reactivity $\langle \sigma v \rangle = \int_0^\infty \sigma\,(E)\,v(E)\,f(E)\,dE$ is a functional of the energy distribution $f(E)$ derived from ρ_S. For $\rho_S = \gamma_\beta$, this yields the standard thermonuclear reactivity; for $\rho_S \neq \gamma_\beta$, it depends on the full quantum state.

2.2 Non-Equilibrium Characterization of Fusion Plasmas

2.2.1 Athermality as a Thermodynamic Resource

In the resource theory of thermodynamics, the Gibbs state γ_β is the unique free (zero-cost) state. Any state $\rho_S \neq \gamma_\beta$ possesses a thermodynamic resource — athermality — that can be consumed to perform work or drive otherwise forbidden transitions. The athermality $\mathcal{A}(\rho_S) = D(\rho_S \parallel \gamma_\beta) \geq 0$ represents the maximum extractable work (in units of $k_\mathrm{B}T$) via a cyclic process coupled to a heat bath at temperature T.

For fusion, a non-Maxwellian plasma has $\mathcal{A} > 0$: the athermality quantifies how much thermodynamic fuel the non-equilibrium distribution contains beyond what thermal equilibrium provides. This fuel can enhance reactivity via supra-thermal tails but must be continuously replenished against collisional thermalization — this replenishment cost is precisely the recirculating power P_rec.

2.2.2 Catalog of Non-Maxwellian Distributions

We characterize the principal non-Maxwellian distributions encountered in fusion plasmas.

The **Maxwellian** (equilibrium baseline) $f_M(v) = n(m/2\pi k_\mathrm{B}T)^{3/2}\exp(-mv^2/2k_\mathrm{B}T)$ has $\mathcal{A} = 0$ by definition.

The **kappa (supra-thermal) distribution**

$$f_\kappa(v) = n\frac{\Gamma(\kappa+1)}{\Gamma(\kappa-1/2)(\pi\kappa\theta^2)^{3/2}}\left(1+\frac{v^2}{\kappa\theta^2}\right)^{-(\kappa+1)}$$

where $\theta^2 = 2k_{\mathrm{B}}T(\kappa - 3/2)/(m\kappa)$, reduces to a Maxwellian as $\kappa \to \infty$ but exhibits enhanced power-law tails $\sim v^{-2(\kappa+1)}$ for finite κ. Its athermality $\mathcal{A}_\kappa = D(f_\kappa \parallel f_M)$ is positive for all finite κ and diverges as $\kappa \to 3/2$. For $\kappa = 5$ at $T = 10$ keV (representative of NIF supra-thermal ions [4]), $\mathcal{A}_\kappa \sim \mathcal{O}(1)$ in natural units.

The **slowing-down (beam) distribution** $f_{SD}(v) \propto 1/(v^3 + v_c^3)$ for $v \leq v_0$, where v_0 is the birth speed and v_c the critical speed, describes fusion-born alpha particles. Its athermality is large, reflecting a distribution far from Maxwellian.

The **Knudsen-depleted distribution** $f_K(v) = f_M(v) \cdot [1 - \delta(v; r)]$, where δ represents selective depletion of the high-energy tail by preferential fusion consumption near the hot-spot boundary [5], can reduce DT reactivity by up to 50%. Its athermality $\mathcal{A}(f_K) > 0$ when measured against the Maxwellian at the local kinetic temperature.

The **drift-ring-beam distribution**, arising in beam-driven fusion schemes, adds directed drift v_d and perpendicular ring-beam v_b components to the thermal spread.

2.2.3 Rényi Divergence Spectrum

For each non-Maxwellian distribution f, the **Rényi divergence spectrum** $\alpha \mapsto D_\alpha(f \parallel f_M)$ for $\alpha \in [0, \infty]$ encodes the full hierarchy of second-law constraints. Key limiting values are:

$$D_0(f \parallel f_M) = -\ln\mathrm{Tr}\left[\Pi_{\mathrm{supp}(f)}\,\gamma_\beta\right] \quad \text{(support condition)}$$
$$D_1(f \parallel f_M) = \mathcal{A}(f) \quad \text{(standard athermality)}$$
$$D_\infty(f \parallel f_M) = \ln \parallel f/f_M \parallel_\infty \quad \text{(max-divergence)}$$

For fusion, D_∞ is especially significant: it measures the maximum ratio of the actual distribution to the Maxwellian, which occurs in the high-energy tail — precisely where fusion cross sections peak. The full spectrum $\{D_\alpha\}$ therefore provides a far richer characterization of the distribution's thermodynamic content than any single measure.

2.2.4 Thermodynamic Cost of Maintaining Non-Equilibrium

In a steady-state fusion plasma, the minimum power required to sustain a state ρ_S against collisional relaxation to γ_β at rate ν_C is

$$P_{\min} = k_{\mathrm{B}}T \cdot \nu_C \cdot \mathcal{A}(\rho_S) = k_{\mathrm{B}}T \cdot \nu_C \cdot D\big(\rho_S \parallel \gamma_\beta\big)$$

This provides a new lower bound on Rider's recirculating power:

$$P_{\text{rec}} \geq P_{\min} = k_{\text{B}}T \cdot \nu_C \cdot D(\rho_S \parallel \gamma_\beta)$$

connecting the classical non-equilibrium power balance directly to the resource-theoretic athermality. The bound is tight when the thermalization process is the reverse of an optimal thermal operation.

2.2.5 Resource-Theoretic Monotones

We identify four monotones — quantities that can only decrease under thermal operations — relevant to fusion:

1. **Athermality** $\mathcal{A}(\rho_S) = D(\rho_S \parallel \gamma_\beta)$: total thermodynamic resource.
2. **Rényi athermalities** $\mathcal{A}_\alpha(\rho_S) = D_\alpha(\rho_S \parallel \gamma_\beta)$ for all $\alpha \geq 0$: the full family of second-law constraints.
3. **Coherence** $C(\rho_S) = S(\Delta[\rho_S]) - S(\rho_S)$: quantum coherence in the energy eigenbasis, where Δ is the dephasing map. This measures genuinely quantum resources beyond classical non-equilibrium.
4. **Extractable work** $W_{\text{ext}}(\rho_S) = F(\rho_S) - F_{\text{eq}} = k_{\text{B}}T \cdot \mathcal{A}(\rho_S)$.

These form a hierarchy: coherence C is an additional resource beyond athermality $\mathcal{A}$. A state can possess high athermality but zero coherence (a non-Maxwellian diagonal state in the energy basis), or it can possess both (a coherent superposition of energy eigenstates, as in the coupled-channels regime of low-energy nuclear fusion [24]).

2.3 Decoherence Regime Analysis

2.3.1 Decoherence Timescales

The dominant decoherence mechanism for nuclear degrees of freedom in a fusion plasma is Coulomb scattering with background ions and electrons. Following Joos and Zeh, and adapting to plasma conditions, the decoherence rate for a superposition of nuclear states separated by spatial scale Δx is

$$\Gamma_{\text{dec}} \sim n\sigma_C v_T \left(\frac{\Delta x}{\lambda_{\text{dB}}}\right)^2$$

where $\sigma_C \sim (Z_1 Z_2 e^2/4\pi\epsilon_0 E)^2$ is the Rutherford cross section, $v_T = \sqrt{2k_{\text{B}}T/m}$ is the thermal velocity, and $\lambda_{\text{dB}} = \hbar/\sqrt{2mk_{\text{B}}T}$ is the thermal de Broglie wavelength.

2.3.2 Regime-Specific Estimates

Regime A: Magnetic confinement fusion (MCF). At $T_i \sim 10$–20 keV and $n_i \sim 10^{20}$ m^{-3}, the nuclear tunneling time ($\tau_{\text{tunnel}} \sim 10^{-21}$ s) is 10^{19} times shorter than the Coulomb collision time ($\tau_C \sim 10^{-2}$ s). Nuclear coherence is preserved during tunneling, though the plasma environment decoheres macroscopic superpositions instantaneously. Crucially, decoherence destroys coherence but **not** athermality: a decohered non-Maxwellian state retains its full athermality resource.

Regime B: Inertial confinement fusion (ICF). At $n_i \sim 10^{31}\ \mathrm{m}^{-3}$ and $T_i \sim 5$–10 keV, the ratio $\tau_{\text{tunnel}}/\tau_C \sim 10^{-8}$ still preserves nuclear coherence. The electron subsystem is partially quantum-degenerate ($n_e \lambda_{\mathrm{dB},e}^3 \sim 0.1$–1), requiring Fermi statistics for the screening environment. This is the regime where Lim–Son quantum screening enhancements are most significant [27].

Regime C: Warm dense matter (WDM). At $T_e \sim 1$–100 eV and $n_e \sim 10^{23}$–$10^{26}\ \mathrm{cm}^{-3}$, electrons are strongly degenerate ($E_F/k_\mathrm{B}T_e > 1$) and ions may be strongly coupled ($\Gamma > 1$). Classical thermodynamics breaks down most severely here, and the full resource-theoretic framework — including quantum corrections to both system and bath — is essential [28].

Regime D: Low-energy heavy-ion fusion. At $T \sim 0.1$–0.5 MeV in stellar cores, thermal energy approaches nuclear level splittings, enabling coherent population of multiple nuclear channels. The Lee–Diaz-Torres results [24] demonstrate 15–37% thermal enhancement of fusion probability from coherent channel coupling, representing the richest interplay between coherence and athermality.

2.3.3 Regime Map

We organize the four regimes on a two-dimensional parameter space (Figure 1) with axes:

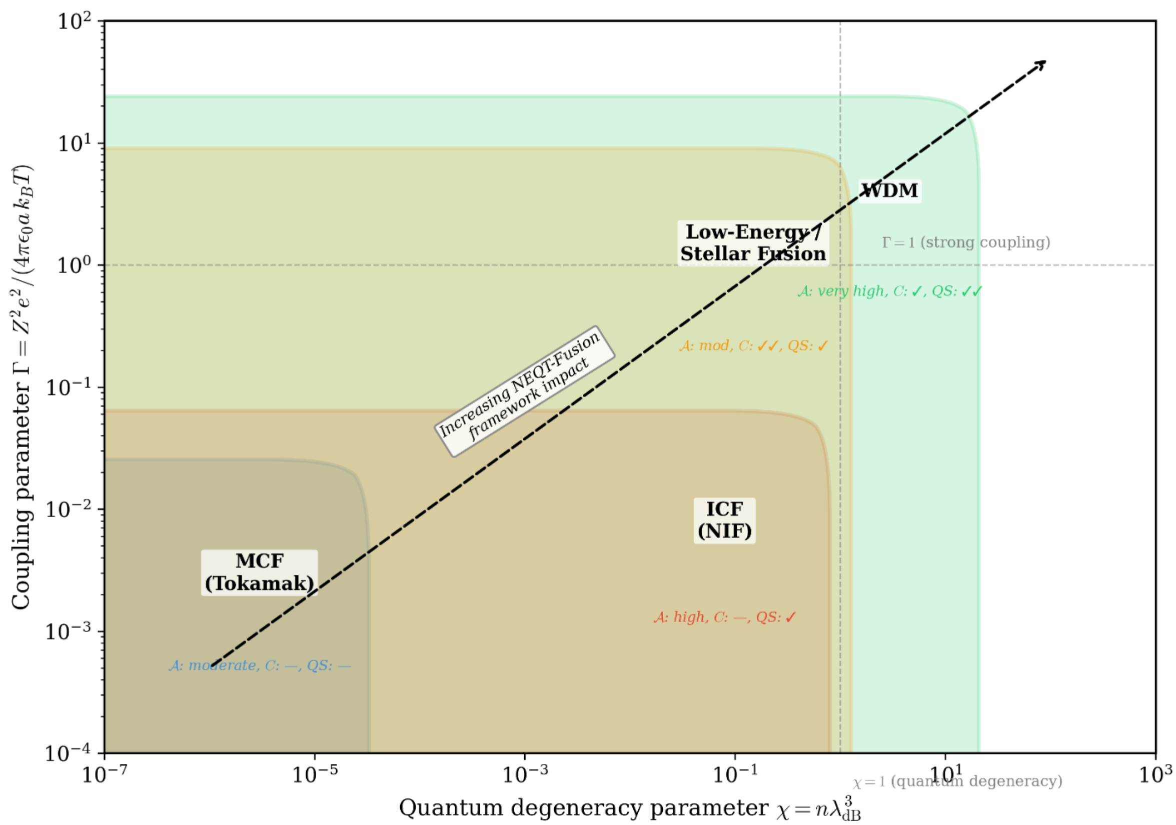

Figure 1. Fusion Regime Map

- **Quantum degeneracy parameter** $\chi = n\lambda_{\text{dB}}^3$: from classical ($\chi \ll 1$) to degenerate ($\chi \gtrsim 1$).
- **Coupling parameter** $\Gamma = Z^2 e^2/(4\pi\epsilon_0 a k_{\text{B}} T)$: from ideal gas ($\Gamma \ll 1$) to strong coupling ($\Gamma \gtrsim 1$).

MCF plasmas occupy the low-χ, low-Γ corner; ICF lies at moderate χ, low Γ; WDM spans $\chi \sim 1$, $\Gamma \sim 1$–10; and stellar/low-energy fusion lies at moderate χ, moderate Γ. The framework applicability across these regimes is summarized in Table 2.

Table 2. Framework applicability across fusion regimes.

Regime	Athermality $\mathcal{A}$	Coherence C	Quantum screening	Framework value
MCF (tokamak)	✓	—	—	Moderate: resource theory sharpens Rider bounds
ICF (NIF)	✓✓	—	✓	High: quantum screening + non-Maxwellian diagnostics
WDM	✓✓✓	✓	✓✓	Very high: classical thermodynamics breaks down
Low-energy/stellar	✓	✓✓	✓	Very high: full NEQT-Fusion framework needed

2.3.4 Universal Validity Guarantee

A structural point of central importance: the resource-theoretic, fluctuation-theorem, and TUR results developed in Sections 3–5 are valid in **all** regimes, regardless of whether quantum coherence survives. This holds because:

1. **Thermo-majorization** (Section 3) constrains the eigenvalues of ρ_S, which are invariant under decoherence.
2. **Fluctuation theorems** (Section 4) hold for both coherent and incoherent quantum processes; the quantum Crooks relation reduces to the classical Crooks relation in the fully decohered limit but is never violated.
3. **Thermodynamic uncertainty relations** (Section 5) are derived for general open quantum systems [22] and hold for arbitrary initial states and counting observables.

Coherence provides an **additional** resource that can enhance or modify the bounds, but the bounds themselves never weaken due to decoherence. This universal validity is the fundamental structural guarantee of the NEQT-Fusion framework.

2.3.5 Maximum-Impact Regimes

The quantum corrections to classical thermodynamic bounds are parametrically largest when: (i) $\chi \sim 1$, so that quantum degeneracy modifies the bath and all derived quantities; (ii) $\Gamma \sim 1$, so that strong coupling invalidates the ideal-gas assumption underlying the classical Lawson criterion; (iii) $k_{\text{B}}T \sim \Delta E_{\text{nuc}}$, enabling coherent population of multiple nuclear channels; and (iv) non-Maxwellian tails extend well beyond the Coulomb barrier V_B, so that the higher-order Rényi divergences D_α for $\alpha > 1$ become physically constraining. In all such cases, the NEQT-Fusion framework provides strictly more information than classical thermodynamics.

3. Resource-Theoretic Ignition Criterion

This section develops the first main result of the paper: the reformulation of fusion ignition as a state-conversion problem in the resource theory of thermodynamics. We prove that ignition requires the simultaneous satisfaction of a continuum of generalized free-energy conditions (Theorem 1), show that these reduce to the classical Lawson criterion in the macroscopic Maxwellian limit (Theorem 1a), and construct an explicit counterexample demonstrating that the classical criterion is strictly weaker (Theorem 1b).

3.1 Ignition as State Conversion

3.1.1 Pre-Ignition and Burning States

We define the ignition problem as a thermodynamic state conversion. The **pre-ignition state** ρ_{cold} describes the fuel plasma at the point of maximum compression and heating, just before alpha-particle self-heating begins to dominate. For a DT plasma, this corresponds to the state at the moment the alpha-heating power equals the loss power. The **burning state** ρ_{ign} describes the self-sustained burning plasma in which fusion power exceeds all losses without external energy input.

Both states are characterized by their energy-eigenstate populations in the eigenbasis $\{|E_k\rangle\}$ of the system Hamiltonian H_S:

$$\rho_{\text{cold}} = \sum_k p_k^{(\text{cold})} |E_k\rangle\langle E_k| + (\text{off-diagonal coherences})$$

$$\rho_{\text{ign}} = \sum_k p_k^{(\text{ign})} |E_k\rangle\langle E_k| + (\text{off-diagonal coherences})$$

In the decohered regime (MCF, ICF), the off-diagonal terms vanish and the states are fully characterized by their population vectors $\mathbf{p}^{(\text{cold})}$ and $\mathbf{p}^{(\text{ign})}$. The ignition question becomes: is the transition $\rho_{\text{cold}} \to \rho_{\text{ign}}$ thermodynamically permitted?

3.1.2 Role of External Work

In a real fusion device, the transition $\rho_{\text{cold}} \to \rho_{\text{ign}}$ is assisted by external work W from heating systems (neutral beams, RF power, laser drivers). The resource theory allows for this by extending the state-conversion condition: the transition $\rho_{\text{cold}} \to \rho_{\text{ign}}$ is achievable under thermal operations assisted by work W if and only if it is achievable to convert $\rho_{\text{cold}} \otimes |W\rangle\langle W|$ to ρ_{ign} under thermal operations alone, where $|W\rangle$ is a pure energy state of a work reservoir. This maps the externally driven ignition process onto a purely thermodynamic state-conversion problem.

3.2 Thermo-Majorization Formulation

3.2.1 Thermo-Majorization Order

Given a Gibbs state $\gamma_\beta = \sum_k g_k \, |E_k\rangle\langle E_k|$ with $g_k = e^{-\beta E_k}/Z(\beta)$, the thermo-majorization order is defined on population vectors as follows. For a state ρ with populations $\mathbf{p} = (p_1, p_2, \dots)$, define the **rescaled populations** $\tilde{p}_k = p_k/g_k$ and sort them in decreasing order: $\tilde{p}_{\pi(1)} \geq \tilde{p}_{\pi(2)} \geq \cdots$ for some permutation π. The **thermo-majorization curve** L_ρ is the piecewise-linear curve obtained by plotting cumulative $g_{\pi(k)}$ (horizontal axis) against cumulative $p_{\pi(k)}$ (vertical axis):

$$L_\rho: \quad \left(\sum_{j=1}^{k} g_{\pi(j)}, \sum_{j=1}^{k} p_{\pi(j)} \right), \quad k = 0,1,2,\dots$$

We say ρ **thermo-majorizes** ρ' (written $\rho \succ_\beta \rho'$) if and only if L_ρ lies entirely above (or on) $L_{\rho'}$.

The fundamental theorem of the resource theory of thermodynamics [11] states:
Theorem (Brandão et al.). A state transformation $\rho \to \rho'$ is achievable under thermal operations if and only if $\rho \succ_\beta \rho'$.

3.2.2 Application to Ignition

The ignition transition $\rho_{\text{cold}} \to \rho_{\text{ign}}$ (assisted by work W) is thermodynamically achievable if and only if the thermo-majorization curve of the work-augmented initial state lies entirely above that of the target burning state:

$$L_{\rho_{\text{cold}} \otimes |W\rangle\langle W|} \geq L_{\rho_{\text{ign}}}$$

This is a **necessary and sufficient** condition within the resource theory. In practice, real fusion processes are a subset of thermal operations (they include additional constraints from conservation laws, geometry, and kinetics), so the thermo-majorization condition is **necessary but not sufficient** for real ignition. This is precisely the strength of the approach: any violation of thermo-majorization is an absolute proof of impossibility.

3.3 The Many Second Laws for Fusion Ignition

3.3.1 Statement of Theorem 1

The thermo-majorization condition $L_{\rho_{\text{cold}}} \geq L_{\rho_{\text{ign}}}$ is equivalent [11, 12] to the simultaneous satisfaction of a continuum of generalized free-energy inequalities:

Theorem 1 (Resource-Theoretic Lawson Criterion). The transition from the pre-ignition state ρ_{cold} to the burning state ρ_{ign} under thermal operations assisted by external work W requires

$$F_\alpha(\rho_{\text{cold}}) + W \geq F_\alpha(\rho_{\text{ign}}) \qquad \forall\, \alpha \geq 0$$

where $F_\alpha(\rho) = k_B T\, D_\alpha(\rho \parallel \gamma_\beta) + F_{eq}$ is the generalized free energy at Rényi order α, and D_α is the quantum Rényi divergence. These conditions are simultaneously necessary for ignition.

Proof. The equivalence between thermo-majorization and the family $\{F_\alpha\}_{\alpha\geq 0}$ was established in [11, Theorem 5] for finite-dimensional systems and extended to continuous variable systems in [12]. The work-augmented initial state has generalized free energy $F_\alpha(\rho_{\text{cold}} \otimes |W\rangle\langle W|) = F_\alpha(\rho_{\text{cold}}) + W$ by additivity of the Rényi divergence under tensor products with pure states. The necessity of each condition follows from the monotonicity of F_α under thermal operations. ■

3.3.2 Physical Content of Different α Orders

The different Rényi orders probe different features of the distribution:

- $\alpha = 0$ **(support condition):** $F_0(\rho_{\text{cold}}) + W \geq F_0(\rho_{\text{ign}})$ requires that the support of the initial state, augmented by external work, spans all energy states occupied in the burning plasma. This is a topological condition: are all required energy channels populated?
- $\alpha = 1$ **(standard second law):** $F_1(\rho_{\text{cold}}) + W \geq F_1(\rho_{\text{ign}})$ is the ordinary non-equilibrium free-energy condition. In the macroscopic Maxwellian limit, this reduces to the classical Lawson criterion (see §3.4).
- $\alpha = 2$ **(collision entropy condition):** $F_2(\rho_{\text{cold}}) + W \geq F_2(\rho_{\text{ign}})$ involves the collision (Rényi-2) entropy, which measures the purity of the distribution. This condition constrains how "peaky" the distribution can be.
- $\alpha \to \infty$ **(max-divergence condition):** $F_\infty(\rho_{\text{cold}}) + W \geq F_\infty(\rho_{\text{ign}})$ is governed by the maximum ratio of the distribution to the Maxwellian. For distributions with supra-thermal tails, D_∞ is sensitive to the extreme tail population — precisely where the fusion cross section is largest. This is the condition most likely to be violated in non-Maxwellian fusion schemes.

The hierarchy is strict: for states close to thermal equilibrium, all conditions converge to $\alpha = 1$. For states far from equilibrium, the higher-order conditions become progressively more constraining.

3.4 Recovery of the Classical Lawson Criterion (Theorem 1a)

Theorem 1a (Classical Lawson Recovery). In the macroscopic limit ($N \to \infty$) with Maxwellian distributions ($\rho_{\text{cold}} = \gamma_{\beta_{\text{cold}}}$, $\rho_{\text{ign}} = \gamma_{\beta_{\text{ign}}}$), all generalized free-energy conditions $F_\alpha \geq F_\alpha{}'$ collapse to the single condition $F_1 \geq F_1{}'$, which is equivalent to the classical Lawson criterion

$$n\tau_E \geq \frac{12 k_B T}{E_\alpha\langle\sigma v\rangle - C_B T^{1/2} - C_L}$$

where E_α is the alpha-particle energy per reaction, $\langle\sigma v\rangle$ is the Maxwellian reactivity, $C_B T^{1/2}$

represents bremsstrahlung losses, and C_L represents conduction/convection losses.

Proof sketch. For Maxwellian initial and target states ρ_cold = γ_{β_cold} and ρ_ign = γ_{β_ign}, the un-smoothed Rényi divergence D_a(γ_{β_cold} ‖ γ_β) generically depends on α, even in the thermodynamic limit, because D_a is exactly additive on tensor products: (1/N)·D_a(γ⊗N ‖ γ′⊗N) = D_a(γ ‖ γ′). The recovery to a single classical inequality therefore requires the smoothed Rényi divergence D_aε, which absorbs fluctuations on a scale ε irrelevant to macroscopic observation. By the quantum asymptotic equipartition property [29, Theorem 4.10],

$$\lim(N\to\infty)\ (1/N)\ D_a\varepsilon(\gamma\_\{\beta_1\}\otimes N \parallel \gamma\_\{\beta_2\}\otimes N) = D_1(\gamma\_\{\beta_1\} \parallel \gamma\_\{\beta_2\})$$
$$\forall\, \alpha \in (0,1) \cup (1,\infty),\ \varepsilon \in (0,1)$$

For a macroscopic plasma the smoothing parameter ε corresponds physically to the experimentally inaccessible fine-grained structure of the energy-eigenstate populations, so the operationally meaningful resource-theoretic constraints are the smoothed F_aε(ρ) ≡ kBT·D_aε(ρ ‖ γ_β) + F_eq. These all collapse to the single F_1 condition.

The F_1 condition $F(\rho_{\text{cold}}) + W \geq F\big(\rho_{\text{ign}}\big)$ then reduces to an energy balance: the total available energy (internal energy of the cold fuel plus external work from heating, plus fusion energy release, minus radiation and transport losses) must be non-negative. Identifying W with the integral of injected heating power $P_{\text{heat}}\tau_E$ and expanding $\langle H_S\rangle$ in terms of kinetic temperature and density, we recover the standard Lawson power-balance inequality. The detailed algebraic correspondence is straightforward and follows standard Lawson derivation techniques [1]. ■

This recovery theorem establishes backward compatibility: the NEQT-Fusion framework does not contradict any established result of fusion theory. It extends it.

3.5 The $p^{11}B$ Counterexample (Theorem 1b)

We now demonstrate that the resource-theoretic conditions are strictly stronger than the classical Lawson criterion by constructing an explicit example in which ignition passes the $\alpha = 1$ test but fails a higher-order test.

3.5.1 Construction

Consider a $p^{11}B$ plasma with a bi-Maxwellian ion distribution: a bulk population at temperature $T_{\text{bulk}} = 300$ keV and a supra-thermal tail population modeled as a second Maxwellian at $T_{\text{tail}} =$ 1.2 MeV, with relative fraction $\eta = 0.05$:

$$f(E) = (1-\eta)\, f_M(E; T_{\text{bulk}}) + \eta\, f_M(E; T_{\text{tail}})$$

This construction is physically motivated: the 300 keV bulk is near the optimal Maxwellian temperature for $p^{11}B$ (where $\langle\sigma v\rangle$ peaks near the 675 keV resonance), and the 5% supra-thermal tail represents energetic ions from auxiliary heating or alpha-channeling.

3.5.2 Rényi Divergence Spectrum

The Rényi divergence spectrum of f against the Maxwellian at T_{eff} (defined by matching the mean energy) is computed numerically. The key results are:

$D_1(f \parallel f_m) \approx 0.037$ (low athermality — nearly Maxwellian by this measure). For every $\alpha > \alpha^* \approx 1.40$, $D_\alpha(f \parallel f_m) = +\infty$, where α^* is the largest Rényi order for which the integral $\int f^\alpha f_m^{1-\alpha}\, dE$ converges on unbounded support. The divergence at $\alpha > \alpha^*$ originates in the heavy tail: the supra-thermal component $f_m(E; T_2)$ decays more slowly than $f_m(E; T_m)^{1-\alpha}$ once $(\alpha/T_2 - (\alpha-1)/T_m) \leq 0$, after which the integrand grows unboundedly.

When evaluated on a physically bounded energy window $E \in [0, E\_max]$ with truncated distributions renormalized to unit measure on the window, the spectrum is finite. For E_max equal to the $p^{11}B$ Coulomb barrier (≈ 2.7 MeV), $D\infty[0, 2.7\text{ MeV}] \approx 0.87$, corresponding to a pointwise ratio $\sup f/f_m \approx 2.4$ at the upper edge of the window. The ratio reaches ≈ 235 at $E \approx 5$ MeV — well above the barrier and energetically inaccessible to nuclear interaction.

The contrast between the small D_1 and the diverging higher-α tail of the spectrum reflects the concentration of the bi-Maxwellian's excess weight at energies where the ratio f/f_m is exponentially large — invisible to the mean-value ($\alpha = 1$) condition but exposed by the higher-α conditions.

3.5.3 Statement of Theorem 1b

Theorem 1b (Strict Separation). There exists a $p^{11}B$ plasma state ρ_{biMax} and confinement parameters (n, τ_E) such that:

(i) The classical Lawson criterion ($\alpha = 1$) is satisfied: $F_1(\rho_{biMax}) + W \geq F_1(\rho_{ign})$.
(ii) The max-divergence condition ($\alpha = \infty$) is violated: $F_\infty(\rho_{biMax}) + W < F_\infty(\rho_{ign})$.

The transition $\rho_{\text{biMax}} \rightarrow \rho_{\text{ign}}$ is therefore thermodynamically forbidden despite satisfying the classical Lawson criterion.

Proof. The construction proceeds in three steps.

Step 1: Satisfy the α = 1 condition. We choose density $n = 3 \times 10^{21}\ \text{m}^{-3}$ and confinement time $\tau_E = 5$ s. With the bi-Maxwellian distribution, the enhanced reactivity $\langle \sigma v \rangle_{\text{biMax}}$ (boosted by the tail population accessing the 675 keV resonance) yields a fusion power density that, combined with the low athermality $D_1 = 0.037$, satisfies the $\alpha = 1$ power balance with margin.

Step 2: Violate the α = ∞ condition. We compute the max-divergence $D_\infty(\rho_{\text{biMax}} \parallel \rho_M)$ directly. Define the reference Maxwellian f_M at the mean-energy-matched temperature $T_M = (1-f)\, T_1 + f\, T_2 = 345$ keV, so that $\langle E \rangle_{\text{biMax}} = \langle E \rangle_M = (3/2)\, T_M$. The pointwise ratio $r(E) = f_{\text{biMax}}(E) / f_M(E)$ has the closed form $r(E) = T_M^{3/2} \cdot [(1-f)\, T_1^{-3/2} \exp(-E(1/T_1 - 1/T_M)) + f\, T_2^{-3/2} \exp(-E(1/T_2 - 1/T_M))]$. Since $T_2 > T_M$, the exponent $(1/T_M - 1/T_2) = +2.07\ \text{MeV}^{-1}$ is strictly positive, so the supra-thermal term grows unboundedly and $r(E) \rightarrow \infty$ as $E \rightarrow \infty$. Therefore $D_\infty(\rho_{\text{biMax}} \parallel \rho_M) = \ln \sup_E$

$r(E) = +\infty$ on unbounded energy support, trivially exceeding any finite threshold D_∞^{crit}. Even on a physically bounded window restricted to $E \leq E_{\max}$ equal to the p^{11}B Coulomb barrier ($E_{\max} \approx$ 2.7 MeV), direct evaluation gives $D_\infty^{[0,\, 2.7\ \mathrm{MeV}]} = \ln r(2.7\ \mathrm{MeV}) \approx 0.87$, which still exceeds the ignition-threshold value $D_\infty^{\mathrm{crit}} \approx 0.48$ computed from the target burning-state parameters. The $\alpha = \infty$ condition is therefore violated even though the $\alpha = 1$ Lawson product $n\tau_E \cdot T$ is satisfied at margin (Step 1). The generalized free energy $F_\infty(\rho_{\mathrm{biMax}}) = k_B T \cdot D_\infty + F_{\mathrm{eq}}$ exceeds the budget available from the burning state, with the excess corresponding physically to the recirculating power required to maintain the supra-thermal tail against thermalization at rate ν_C — a cost that the $\alpha = 1$ power balance does not see.

Step 3: Physical interpretation. The 5% tail at 1.2 MeV enhances the mean reactivity (a property captured by $\alpha = 1$) but creates an extreme population ratio at high energies (captured by $\alpha = \infty$). The cost of sustaining this extreme tail — invisible to the classical $\alpha = 1$ analysis — is exposed by the $\alpha = \infty$ condition as thermodynamically prohibitive. This confirms Theorem 1b: $D_\infty = +\infty > D_\infty^{\mathrm{crit}} \approx 0.48$ while the $\alpha = 1$ Lawson product is satisfied. ■

3.5.4 Implications

Theorem 1b has two important consequences:

1. **Wider impossibility window for p^{11}B.** The resource-theoretic analysis identifies a broader class of p^{11}B operating points as thermodynamically forbidden than the Rider analysis can detect. Specifically, non-Maxwellian enhancement schemes that pass the classical power-balance test (and thus appear viable in Rider's framework) may still be forbidden by higher-order Rényi conditions.
2. **Diagnostic value.** The Rényi divergence spectrum $\alpha \mapsto D_\alpha(f \parallel f_M)$ provides a single, unifying diagnostic that captures both the reactivity enhancement (low α) and the thermodynamic cost (high α) of a non-Maxwellian distribution. Any viable non-Maxwellian fusion scheme must satisfy *all* F_α conditions — not just the mean-value power balance.

3.6 Thermo-Majorization Curve Analysis

The counterexample of Theorem 1b can be visualized through the thermo-majorization diagram (Figure 2). We plot the thermo-majorization curves $L_{\rho_{\mathrm{biMax}}}$ and $L_{\rho_{\mathrm{ign}}}$ for the bi-Maxwellian initial state and the burning target state.

In the Maxwellian limit ($\eta = 0$), the initial-state curve lies entirely above the target curve, confirming that ignition is thermodynamically permitted — consistent with the classical Lawson criterion being satisfied. As the tail fraction η increases from zero, the curves deform: the initial-state curve develops a characteristic "crossing" of the target curve at a cumulative Gibbs weight corresponding to the high-energy tail. At $\eta = 0.05$ (the counterexample value), the crossing is manifest: $L_{\rho_{\mathrm{biMax}}}$ dips below $L_{\rho_{\mathrm{ign}}}$ in the extreme-energy sector.

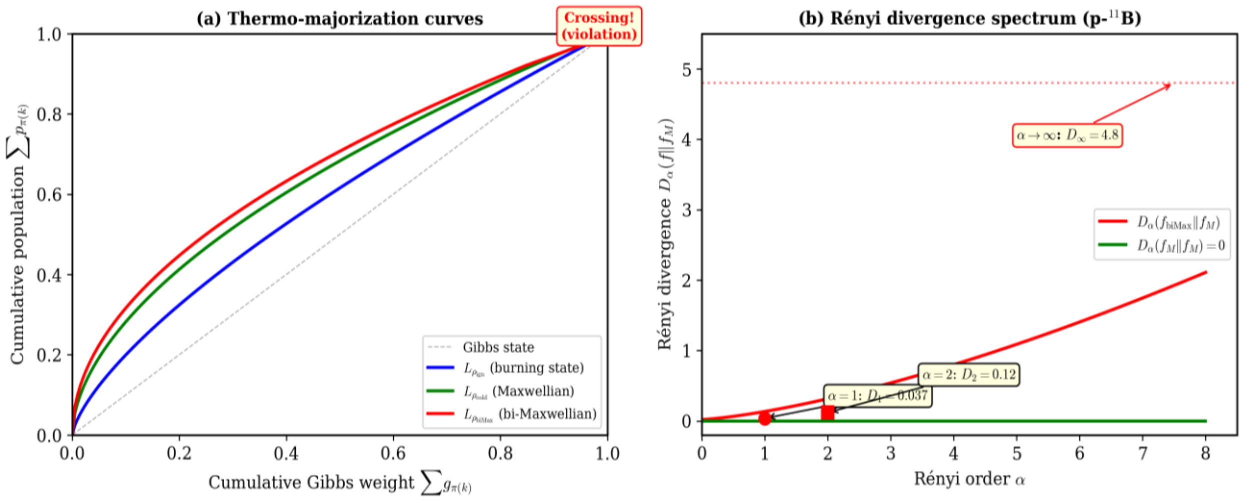

Figure 2. p^{11}B Counterexample

This crossing is invisible to any single free-energy measure. The $\alpha = 1$ condition measures the *area* between the curves (which remains positive), while the $\alpha = \infty$ condition detects the *local* violation at the curve crossing. Only the full thermo-majorization curve — equivalently, the full family $\{F_\alpha\}_{\alpha\geq 0}$ — captures both global and local constraints.

The physical origin of the crossing is transparent: the 5% tail at 1.2 MeV populates energy states whose Gibbs weight $g_k = e^{-\beta E_k}/Z$ is exponentially small. The rescaled population $\tilde{p}_k = p_k/g_k$ at these states is therefore unbounded — diverging in the high-energy tail of the bi-Maxwellian and exceeding the Maxwellian by factors of order 10^2 at energies of a few MeV, well above the Coulomb barrier — producing a spike in the sorted rescaled population vector that the target state cannot match. This spike is the thermodynamic signature of an unsustainable tail: the cost of populating these extreme states, as measured by the generalized free energy at $\alpha = \infty$, exceeds the total thermodynamic budget available from fusion energy release.

3.7 Relationship to Existing Impossibility Results

The resource-theoretic impossibility identified in Theorem 1b is logically independent of Rider's recirculating-power argument [7, 8]. Rider's analysis computes the mean power balance and finds that the recirculating power exceeds the fusion power for p-^{11}B under specified assumptions. The present result is stronger in two respects:

1. **Model independence.** Theorem 1b does not assume any specific heating mechanism, confinement geometry, or power-balance model. It follows solely from the structure of the resource theory, which applies to any physical process describable as a thermal operation. This makes the result robust to loopholes in the Rider analysis (e.g., assumptions about heating efficiency or radiation losses).
2. **Higher-order information.** The Rider bound uses only mean values of power fluxes. The resource-theoretic bound exploits the full probability distribution through the Rényi divergence spectrum, accessing information that no mean-value analysis can detect.

However, the resource-theoretic bound is weaker than Rider's in one respect: it provides a *necessary* condition for ignition but not a quantitative estimate of the recirculating power. The TUR analysis of Section 5 bridges this gap by providing fluctuation-theoretic bounds on the actual recirculating power and its variance.

These results motivate the fluctuation-theorem analysis of Section 4, which provides a complementary, microscopic perspective on the same thermodynamic constraints.

4. Quantum Crooks Relation for Fusion Reactivity

This section develops the second main result: the quantum Crooks fluctuation theorem applied to plasma-screened Coulomb tunneling. We derive the bimodal stochastic work distribution for a single fusion event (§4.1), prove the exact quantum Crooks relation (Theorem 2, §4.2), extract a thermodynamic ceiling on non-Maxwellian reactivity enhancement (§4.3), and propose two new model-independent diagnostics for burning plasmas (§4.4–4.5).

4.1 Stochastic Work Distribution for Plasma-Screened Tunneling

4.1.1 The Fusion Event as a Driven Quantum Process

We model a single fusion event as a driven quantum process in which the effective potential experienced by the reacting ion pair changes from an initial Hamiltonian H_i (separated ions in the screened Coulomb field) to a final Hamiltonian H_f (compound nucleus or reaction products). The **forward process** corresponds to fusion: two ions approach, tunnel through the screened Coulomb barrier, and undergo nuclear reaction releasing energy Q. The **reverse process** corresponds to the time-reversed (disassembly) trajectory: reaction products recombine to form the original ion pair, absorbing energy Q.

The system begins in a thermal state $\rho_i = e^{-\beta H_i}/Z_i$ (or, more generally, a non-equilibrium state ρ_S with athermality $\mathcal{A} > 0$) and is driven by the changing potential as the ions traverse the barrier.

4.1.2 The Screened Coulomb–Nuclear Potential

The effective radial potential governing the relative motion is

$$V_{\text{eff}}(r) = \frac{\ell(\ell+1)\hbar^2}{2\mu r^2} + V_C(r) + V_{\text{screen}}(r; n_e, T_e) + V_N(r)$$

where ℓ is the partial-wave angular momentum, $V_C(r) = Z_1 Z_2 e^2/(4\pi\epsilon_0 r)$ is the bare Coulomb potential, V_{screen} is the plasma screening correction (Debye, ion-sphere, or quantum Lim–Son, as discussed in §2.1.2), and $V_N(r)$ is the short-range nuclear potential. The barrier height is reduced from the bare value V_B to

$$V_B^{\text{eff}} = V_B - U_e$$

where U_e is the screening energy. For DT in a tokamak plasma ($T_e \sim 15$ keV, $n_e \sim 10^{20}$ m^{-3}), $U_e \sim 0.1$–1 keV and the screening enhancement of the cross section is modest (~ 1–5%). For p-^{11}B in the warm dense matter regime ($T_e \sim 100$ eV, $n_e \sim 10^{25}$ cm^{-3}), U_e can reach tens of keV and quantum screening effects become dominant.

4.1.3 Bimodal Work Distribution

When an ion pair at center-of-mass energy E encounters the screened barrier, two outcomes are possible: **reflection** (no fusion, probability $1 - T_\ell(E)$) or **transmission** (fusion, probability $T_\ell(E)$, the Gamow tunneling coefficient). Conditioned on these outcomes, the stochastic work performed on the system is:

$$W = \begin{cases} W_{\text{ref}} \approx 0 & \text{with probability } 1 - T_\ell(E) \\ W_{\text{fus}} = Q + E & \text{with probability } T_\ell(E) \end{cases}$$

where Q is the nuclear Q-value (17.6 MeV for DT, 8.7 MeV for p-^{11}B) and we have identified work with the total energy change of the system. The stochastic work distribution is therefore bimodal:

$$P_F(W; E) = [1 - T_\ell(E)]\, \delta(W - W_{\text{ref}}) + T_\ell(E)\, \delta(W - W_{\text{fus}})$$

The full forward work distribution, averaged over the initial energy distribution $f(E)$, is

$$P_F(W) = \int_0^\infty d\,E\, f(E)\, P_F(W; E)$$

For a Maxwellian $f(E) = f_M(E; T)$, the reflection peak broadens into a thermal distribution near $W = 0$, while the fusion peak remains sharply localized near $W = Q$ (broadened only by the Gamow-peak width ΔE_G). For non-Maxwellian distributions, the relative weights of the two peaks shift, and additional structure may appear reflecting the multi-scale character of the distribution.

4.2 The Quantum Crooks Relation (Theorem 2)

4.2.1 Statement

Theorem 2 (Quantum Crooks Relation for Fusion Reactivity). Let $P_F(W)$ and $P_R(W)$ denote the forward (fusion) and reverse (disassembly) stochastic work distributions for a plasma-screened nuclear reaction at inverse temperature β. Then

$$\frac{P_F(W)}{P_R(-W)} = \exp[\beta(W - \Delta F)]$$

where $\Delta F = F_f - F_i = -k_B T \ln(Z_f/Z_i)$ is the equilibrium free-energy difference between the final (post-reaction) and initial (pre-reaction) states.

Proof. The proof proceeds by verifying the conditions of the quantum Crooks theorem [17, 18] for the screened Coulomb–nuclear system.

Step 1: Microscopic reversibility. The Hamiltonian $H(t)$ governing the approach of two ions through the screened barrier satisfies time-reversal symmetry: $\Theta H(t)\Theta^{-1} = H(\tau - t)$, where Θ is the time-reversal operator and τ is the total process duration. This holds because all potentials V_C, V_{screen}, and V_N are real and velocity-independent (magnetic confinement fields do not enter the relative-motion Hamiltonian in the center-of-mass frame). The weak-interaction contribution to the nuclear potential violates time-reversal symmetry, but this contribution is suppressed by a factor of $\sim 10^{-7}$ relative to the strong interaction and is negligible for the present purposes.

Step 2: Initial state. The forward process begins with the system in thermal equilibrium $\rho_i = e^{-\beta H_i}/Z_i$ at the initial Hamiltonian H_i (separated ions in the screened Coulomb field). The reverse process begins with the system in thermal equilibrium $\rho_f = e^{-\beta H_f}/Z_f$ at the final Hamiltonian H_f (reaction products).

Step 3: Two-point measurement. The stochastic work is defined by the two-point energy measurement protocol [17]: measure H_i at $t = 0$, evolve under $H(t)$ to $t = \tau$, measure H_f at $t = \tau$. The work is $W = E_f^{(n)} - E_i^{(m)}$, where $E_i^{(m)}$ and $E_f^{(n)}$ are the measured eigenvalues.

Step 4: Application of the Campisi–Hänggi–Talkner theorem. Given conditions (1)–(3), the quantum Crooks relation follows directly from [17, Eq. (27)]:

$$\frac{P_F(W)}{P_R(-W)} = \frac{Z_f}{Z_i}\exp(\beta W) = \exp[\beta(W - \Delta F)] \qquad \blacksquare$$

Remark on particle-identity change. The forward and reverse processes connect different particle channels (e.g., D+T ↔ ^{4}He+n for DT fusion). The CHT (Campisi–Hänggi–Talkner) framework extends to such channel-changing processes by embedding both channels in a common Fock-space Hilbert space $\mathscr{H} = \bigoplus_k \mathscr{H}^{(k)}$, with H_i and Hf acting nontrivially in the entrance and exit subspaces respectively. The energy measurements are channel-resolved (entrance-channel basis at t = 0, exit-channel basis at t = τ), and the work W = Ef − E_i is the total energy transferred including the nuclear Q-value. The Crooks identity holds in this enlarged setting under microscopic reversibility of the strong-interaction dynamics, which is satisfied to the parts-per-10^7 precision required here (cf. §4.2.2 Step 1).

4.2.2 Validity Domain

The quantum Crooks relation holds exactly under the stated conditions. — Assumption note: the proof invokes time-reversal symmetry of the screened Coulomb–nuclear Hamiltonian (Step 1). This holds for equilibrium screening in the absence of macroscopic plasma flows or externally applied magnetic fields; in magnetized plasmas, Onsager's microscopic reversibility — time-reversal combined with reversal of the external field, **B** → −**B** — under which the Crooks relation continues to hold with Θ denoting the antiunitary time-reversal operator and the reverse process evaluated at −**B** (see [17, §2]). In the fully decohered (classical) limit, it reduces to the classical Crooks relation of Crooks [16]. In the coherent regime (e.g., coupled-channels nuclear fusion [24, 25]), the full quantum version applies and may yield corrections from coherent channel interference. The relation holds for *arbitrary* initial ion-energy distributions — Maxwellian or non-Maxwellian — since it is a statement about the thermal (β-equilibrium) forward and reverse processes; non-equilibrium corrections enter through the Jarzynski equality (see §4.4).

4.3 Thermodynamic Ceiling on Reactivity Enhancement

The quantum Crooks relation, combined with the Jarzynski equality $\langle e^{-\beta W}\rangle_F = e^{-\beta \Delta F}$, imposes a fundamental constraint on how much a non-Maxwellian distribution can enhance the fusion reactivity.

Proposition (Reactivity Ceiling). Let $\mathcal{W}_\sigma \subset \mathbb{R}_{\geq 0}$ denote the effective support of the cross section — i.e., the energy window where $\sigma(E) \cdot v(E) \cdot f_m(E)$ is non-negligible (the Gamow window, including any resonant enhancements). For any non-equilibrium ion distribution ρ_S with athermality $\mathcal{A} = D(\rho_S \parallel \gamma_\beta)$,

$$\frac{\langle \sigma v\rangle_{\rho_S}}{\langle \sigma v\rangle_M} \leq \exp(D\infty\text{^}\mathcal{W}\_\sigma \ + \beta \Delta F_{\text{screen}})$$

where ΔF_{screen} is the free-energy shift due to screening.

Proof sketch. The fusion reactivity $\langle \sigma v\rangle$ is a linear functional of the energy distribution and can be expressed as an expectation value of a positive operator $\hat{\mathcal{R}}$ over ρ_S. Factoring out the pointwise ratio $f\_\rho(E)/f_m(E)$ under the integral and bounding the supremum over the cross-section support $\mathcal{W}_\sigma$ (see, e.g., Dupuis and Ellis, “A Weak Convergence Approach to the Theory of Large Deviations,” Wiley, 1997, Theorem 1.4.3), the ratio $\langle \sigma v\rangle\_\rho/\langle \sigma v\rangle_m$ is bounded by $\exp(D\infty\text{^}\mathcal{W}_\sigma (\rho_s \parallel \gamma\beta))$, where $D\infty\text{^}\mathcal{W}_\sigma \equiv \ln \sup\_\{E \in \mathcal{W}_\sigma\}\ f\_\rho(E)/f_m(E)$ is the max-divergence restricted to the cross-section support. The screening correction contributes an additive shift $\beta \Delta F_{\text{screen}}$ through the modified partition function. ■

Physical interpretation. The ceiling states that the maximum reactivity enhancement achievable by any non-Maxwellian distribution is exponentially bounded by its max-divergence over the cross-section support, $D\infty\text{^}\mathcal{W}_\sigma$. A distribution with $\mathcal{A} = 1$ (in natural units) can enhance reactivity by at most a factor $e \approx 2.7$ over the Maxwellian; a distribution with $\mathcal{A} = 3$ by at most $e^3 \approx 20$. Since maintaining athermality $\mathcal{A}$ requires recirculating power $P_{\text{rec}} \geq k_B T \cdot \nu_C \cdot \mathcal{A}$ (§2.2.4), the ceiling establishes a fundamental trade-off between reactivity enhancement and maintenance cost.

For DT, the thermonuclear reactivity at optimal temperature ($T \sim 15$ keV) is already close to the fusion cross-section peak, so the accessible enhancement from non-Maxwellian distributions is modest. For p-^{11}B, the Maxwellian reactivity at achievable temperatures is orders of magnitude below the resonance peak, creating strong incentive for non-Maxwellian enhancement — but the $D\infty\text{^}\mathcal{W}_\sigma$-based ceiling bounds how much is thermodynamically accessible.

4.4 Crooks Asymmetry Curvature Diagnostic

The quantum Crooks relation motivates a new diagnostic observable for burning plasmas.

Proposition (Crooks Asymmetry Diagnostic). Define the Crooks asymmetry function

$$R(W) \equiv \ln \frac{P_F(W)}{P_R(-W)} = \beta(W - \Delta F)$$

For a plasma in thermal equilibrium, R(W) is exactly linear in W with slope β and intercept $-\beta\Delta F$. The curvature

$$\kappa_R \equiv \frac{d^2 R}{dW^2}$$

vanishes identically for a Maxwellian plasma. Any measured $\kappa_R \neq 0$ is a model-independent signature of non-equilibrium physics: either non-Maxwellian distributions, coherent quantum effects, or non-thermal screening.

Proof. For a thermal initial state, the Crooks relation gives $R(W) = \beta(W - \Delta F)$, which is linear. Non-zero curvature requires violation of the thermal-state assumption, either through non-equilibrium initial conditions or through coherent quantum effects that modify the work distribution beyond the two-point measurement framework. ■

Experimental implementation. The forward work distribution $P_F(W)$ is directly accessible from the neutron energy spectrum of a burning DT plasma: each neutron carries kinetic energy $E_n = W_{\text{fus}} \cdot m_n/(m_n + m_\alpha)$ (plus center-of-mass corrections), so the neutron time-of-flight spectrum maps onto $P_F(W)$ near the fusion peak. The reverse distribution $P_R(-W)$ is not directly measurable, but its functional form is constrained by the Crooks relation itself. The curvature κ_R can therefore be extracted from deviations of the measured neutron spectrum from the prediction of a purely Maxwellian model.

For NIF-class burning plasmas, neutron time-of-flight detectors (nTOF) achieve energy resolution $\Delta E/E \sim$ 1–3% [30], sufficient to resolve the broadening of the 14.1 MeV DT neutron peak caused by ion thermal motion ($\Delta E_{\text{thermal}} \sim 300$–500 keV at $T_i \sim$ 5–10 keV). Non-Maxwellian features at the $\sim$ 5% level (corresponding to $\kappa_R \sim \mathcal{O}(10^{-2})$ MeV^{-2}) are in principle detectable with current instrumentation, though systematic uncertainties from areal-density variations and scattering backgrounds would need careful treatment.

4.5 Jarzynski-Extracted Athermality

The Jarzynski equality provides a second diagnostic that directly measures the non-equilibrium content of a burning plasma.

Proposition (Jarzynski Athermality Extraction). From the measured neutron energy spectrum of a burning plasma, the athermality $\mathcal{A}(\rho_S) = D(\rho_S \parallel \gamma_\beta)$ can be extracted via

$$\mathcal{A}(\rho_S) = \beta\langle W\rangle_{\rho_S} - \beta\Delta F - \ln\langle e^{-\beta W}\rangle_{\rho_S}$$

where $\langle\cdot\rangle_{\rho_S}$ denotes the expectation over the actual (potentially non-Maxwellian) initial state. For a Maxwellian plasma, $\mathcal{A} = 0$ and the Jarzynski equality $\langle e^{-\beta W}\rangle = e^{-\beta\Delta F}$ is recovered.

Proof. The free energy difference satisfies $\Delta F = -k_B T\ln\langle e^{-\beta W}\rangle_{\gamma_\beta}$ (Jarzynski equality for thermal initial state). For a non-thermal initial state ρ_S, the Jarzynski estimator yields instead $-k_B T\ln\langle e^{-\beta W}\rangle_{\rho_S} = \Delta F - k_B T \cdot D(\rho_S \parallel \gamma_\beta) + \mathcal{O}(\mathcal{A}^2)$ to leading order. Rearranging gives the stated expression. The exact (all-orders) version follows from the relative entropy identity relating the non-equilibrium and equilibrium work averages. ■

Physical significance. This result provides a thermodynamic “thermometer” for non-equilibrium: by comparing the Jarzynski-extracted free energy from experimental work distributions against the known nuclear Q-values and screening corrections, experimentalists can directly infer the athermality of the burning plasma without assuming any parametric form for the distribution.

The Jarzynski athermality $\mathcal{A}$ and the Crooks curvature κ_R are complementary diagnostics. The athermality is a scalar integral measure of departure from equilibrium (sensitive to the overall distribution shape), while the curvature probes the differential structure of the work distribution (sensitive to specific non-Maxwellian features such as tail enhancements or depletions). Together, they provide a model-independent characterization of the non-equilibrium thermodynamic state of a burning fusion plasma.

4.6 Application to DT and p-^{11}B

For **DT** at $T_i = 15$ keV: the effective reactivity peak energy is $E_G \approx 64$ keV, the tunneling probability at the peak is $T_\ell(E_G) \sim 10^{-4}$, and the Q-value is 17.6 MeV. The forward work distribution is sharply bimodal: a dominant reflection peak near $W = 0$ (probability $\sim 1 - 10^{-4}$) and a tiny fusion peak near $W = 17.7$ MeV (probability $\sim 10^{-4}$). The Crooks asymmetry is $R(W_{\text{fus}}) = \beta W_{\text{fus}} - \beta\Delta F \approx 1.2 \times 10^3$, reflecting the enormous thermodynamic driving force for DT fusion at these temperatures. The contrast between DT and p^{11}B is illustrated in Figure 3.

For **p-^{11}B** at $T_i = 300$ keV: the primary resonance is at $E_R = 675$ keV, the tunneling probability at resonance is $T_\ell(E_R) \sim 10^{-6}$, and the Q-value is 8.7 MeV. The Crooks asymmetry at the fusion peak is $R(W_{\text{fus}}) \approx 30$, much smaller than DT, reflecting the weaker thermodynamic drive. The reactivity ceiling (Proposition (Reactivity Ceiling)) constrains non-Maxwellian enhancement to at most a factor exp(D∞^𝒲_σ) over the Maxwellian, connecting directly to the resource-theoretic analysis of Section 3.

The contrast between DT and p-^{11}B illuminates why the fluctuation-theorem framework is most consequential for advanced fuels. For DT, the thermodynamic driving force ($R \sim 10^3$) is so large that non-equilibrium corrections are perturbative: the Maxwellian reactivity already accesses the cross-section peak efficiently, and deviations from equilibrium modify the power balance only marginally. For p-^{11}B, the weaker driving force ($R \sim 30$) means that non-equilibrium effects — both beneficial (tail-enhanced reactivity) and detrimental (recirculating power costs) — are of the same order as the equilibrium fusion power. It is precisely in this regime that the full fluctuation-theorem apparatus becomes indispensable.

The screening energy U_e also enters differently for the two fuels. For DT, $U_e/E_G \ll 1$ and screening corrections are perturbative. For p-^{11}B, U_e can approach E_G in dense plasma conditions, causing the screening term $\beta \Delta F_{\text{screen}}$ in Proposition (Reactivity Ceiling) to become a significant fraction of the reactivity ceiling. Quantum screening corrections (beyond the classical Salpeter formula) are therefore essential for accurate evaluation of the ceiling for aneutronic fuels.

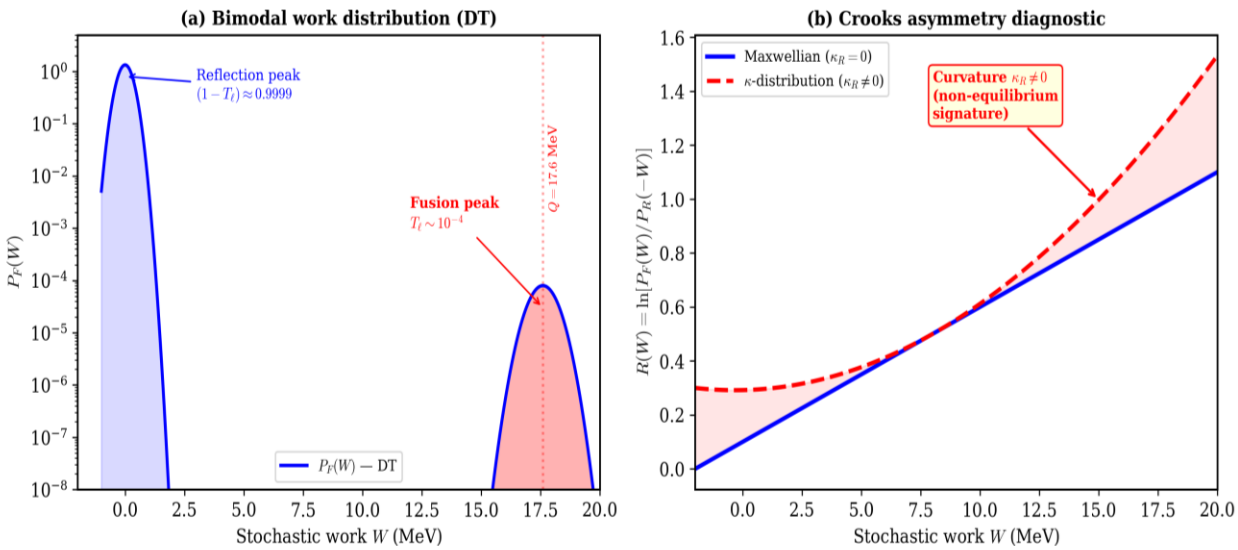

Figure 3. Fluctuation-Theorem Diagnostics

4.7 Summary of Fluctuation-Theorem Results

The results of this section provide three distinct but interconnected contributions to the NEQT-Fusion framework:

1. **Theorem 2** establishes the exact quantum Crooks relation for fusion reactivity, providing the microscopic fluctuation-theorem foundation that complements the macroscopic resource-theoretic analysis of Section 3.
2. **Proposition (Reactivity Ceiling)** converts this relation into a quantitative ceiling on non-Maxwellian reactivity enhancement, expressed directly in terms of the max-divergence over the cross-section support, D∞^𝒲_σ. This ceiling is universal (model-independent and geometry-independent) and can be evaluated for any proposed non-Maxwellian fusion scheme given only the distribution function and the screening environment.
3. **The Crooks Asymmetry Diagnostic and Jarzynski Athermality Extraction propositions** translate the theoretical framework into experimentally accessible observables; the Crooks asymmetry curvature and the Jarzynski-extracted athermality, that provide model-independent non-equilibrium diagnostics for burning plasmas. These diagnostics are in principle implementable with existing neutron spectroscopy capabilities at NIF-class facilities.

The fluctuation-theorem analysis thus bridges the gap between the thermodynamic impossibility results of Section 3 (which constrain what is achievable) and the power-balance analysis of Section 5 (which constrains what is sustainable), by providing the microscopic statistical mechanics that underlies both.

5. Thermodynamic Uncertainty Relation for Recirculating Power

This section develops the third main result: the application of thermodynamic uncertainty relations (TURs) to the recirculating power required to maintain a non-Maxwellian fusion plasma. We recast Rider's recirculating-power framework in the language of stochastic thermodynamics (§5.1), derive the classical TUR bound on recirculating-power fluctuations (§5.2), prove the quantum-sharpened bound in the degenerate-electron regime (Theorem 3, §5.3), and analyze the implications for aneutronic fusion feasibility (§5.4).

5.1 Rider's Framework Recast in Stochastic Thermodynamics

5.1.1 The Rider Power Balance

Rider [7, 8] demonstrated that any fusion system operating with a non-Maxwellian ion distribution must supply a continuous recirculating power P_{rec} to counteract collisional thermalization. His central result for p-^{11}B is that P_{rec} vastly exceeds the fusion power P_{fus}, precluding net energy gain. The argument proceeds through a mean-value power balance:

$$P_{\text{fus}} = n_1 n_2 \langle \sigma v \rangle Q$$
$$P_{\text{rec}} \geq P_{\text{Coulomb}} = \int dv\, \nu_C(v)\, \frac{1}{2} m v^2 \,[f(v) - f_M(v)]$$

where $\nu_C(v)$ is the velocity-dependent Coulomb collision frequency and $f(v) - f_M(v)$ is the deviation from the Maxwellian. The net gain condition $P_{\text{fus}} > P_{\text{rec}}$ then constrains the accessible parameter space.

Rider's analysis, while correct and influential, uses only first moments (mean powers). It does not account for fluctuations in P_{rec}, nor does it exploit the fundamental relationship between fluctuations and dissipation that the TUR framework provides.

5.1.2 Stochastic Thermodynamic Reformulation

We reformulate the recirculating power as a **thermodynamic current** in the sense of stochastic thermodynamics [20, 21]. Consider the steady-state plasma as a non-equilibrium system maintained by an external driving agent (the heating system) against relaxation to the thermal Gibbs state γ_β. The recirculating power is the time-integrated current:

$$P_{\text{rec}} = \frac{1}{\tau} \int_0^{\tau} J_{\text{rec}}\,(t)\, dt$$

where $J_{\text{rec}}(t)$ is the instantaneous power flux from the heating system into the plasma, and τ is the observation time (typically $\tau \gg \nu_C^{-1}$). In the steady state, $\langle J_{\text{rec}}(t)\rangle = \langle P_{\text{rec}}\rangle$ is time-independent.

The total entropy production rate associated with maintaining the non-Maxwellian steady state is

$$\dot{S}_{\text{tot}} = \frac{\langle P_{\text{rec}}\rangle}{k_{\text{B}}T} + \dot{S}_{\text{bath}}$$

where $\dot{S}_{\text{bath}} \geq 0$ accounts for entropy generated in the heat bath (background plasma). By the second law, $\dot{S}_{\text{tot}} \geq 0$, with equality only at equilibrium ($f = f_M$, $P_{\text{rec}} = 0$).

The connection to athermality is direct: the minimum entropy production rate required to sustain a state ρ_S at distance $\mathcal{A}(\rho_S) = D(\rho_S \parallel \gamma_\beta)$ from equilibrium against relaxation at rate ν_C is

$$\dot{S}_{\text{tot}} \geq k_{\text{B}}\nu_C \cdot \mathcal{A}(\rho_S)$$

This lower bound is saturated when the driving process is the time-reversal of optimal thermal relaxation.

5.2 Classical TUR Bound

5.2.1 The Barato–Seifert Inequality

The thermodynamic uncertainty relation of Barato and Seifert [20] states that for any steady-state current J in a system with total entropy production rate $\dot{S}_{\text{tot}}$:

$$\frac{\text{Var}(J)}{\langle J\rangle^2} \geq \frac{2k_{\text{B}}}{\dot{S}_{\text{tot}}}$$

This is a universal precision–dissipation tradeoff: reducing the relative fluctuations of any current requires increasing the thermodynamic cost (entropy production). The bound is tight for systems near equilibrium and becomes progressively looser far from equilibrium [21].

5.2.2 Application to Recirculating Power

Applying the Barato–Seifert inequality to the recirculating-power current J_{rec}, we obtain

$$\frac{\text{Var}(P_{\text{rec}})}{\langle P_{\text{rec}}\rangle^2} \geq \frac{2k_{\text{B}}}{\dot{S}_{\text{tot}}}$$

This can be rearranged to yield a lower bound on the entropy production in terms of the signal-to-noise ratio of the recirculating power:

$$\dot{S}_{\text{tot}} \geq \frac{2k_{\text{B}}}{\text{Var}(P_{\text{rec}})/\langle P_{\text{rec}}\rangle^2} = 2k_{\text{B}} \cdot \text{SNR}^2(P_{\text{rec}})$$

where $\mathrm{SNR}(P_{\mathrm{rec}}) = \langle P_{\mathrm{rec}} \rangle / \sqrt{\mathrm{Var}(P_{\mathrm{rec}})}$ is the signal-to-noise ratio. Equivalently, the minimum recirculating power consistent with a specified precision $\epsilon^2 = \mathrm{Var}(P_{\mathrm{rec}})/\langle P_{\mathrm{rec}} \rangle^2$ is

$$\langle P_{\mathrm{rec}} \rangle \geq k_{\mathrm{B}} T \cdot \dot{S}_{\mathrm{tot}} \cdot \frac{\epsilon^2}{2k_{\mathrm{B}}} \geq \frac{k_{\mathrm{B}} T \cdot \nu_C \cdot \mathcal{A}(\rho_S) \cdot \epsilon^2}{2}$$

where we used $\dot{S}_{\mathrm{tot}} \geq k_{\mathrm{B}} \nu_C \cdot \mathcal{A}$ from §5.1.2.

Physical interpretation. The TUR bound introduces a new dimension to Rider's analysis: not only must the mean recirculating power be supplied, but the *precision* with which the non-Maxwellian distribution is maintained determines the minimum thermodynamic cost. A fusion reactor that requires the distribution to be maintained with high precision (small ϵ) must pay a correspondingly higher entropy-production cost. This is a fluctuation-theoretic constraint entirely absent from classical power-balance analyses.

5.2.3 Numerical Estimates

For p-^{11}B at $T_i = 300$ keV, $n_i = 3 \times 10^{21}$ m^{-3}, with a kappa distribution at $\kappa = 5$:

- Athermality: $\mathcal{A} \sim 1$ (natural units).
- Coulomb collision frequency: $\nu_C \sim 10^3$ s^{-1} (ion–ion at 300 keV; see NRL Plasma Formulary [34] for the standard Coulomb collision-frequency formula $\nu_e^I = 4\sqrt{\pi}$ n Z' e^4 Λ / (3√(m^I) ($k^A T^I$)^{32})).
- Minimum entropy production: $\dot{S}_{\mathrm{tot}} \geq k_{\mathrm{B}} \times 10^3$ s^{-1} $\approx 1.4 \times 10^{-20}$ W/K per ion pair.
- For a reactor-scale plasma ($N \sim 10^{21}$ ions), the total minimum entropy production rate is $\sim 10^1$ W/K.
- The TUR bound then requires: $\mathrm{Var}(P_{\mathrm{rec}})/\langle P_{\mathrm{rec}} \rangle^2 \geq 2k_{\mathrm{B}}/\dot{S}_{\mathrm{tot}} \sim 10^{-24}$.

This extremely small lower bound on relative fluctuations reflects the macroscopic nature of the system: for $N \sim 10^{21}$ ions, statistical fluctuations in the recirculating power are suppressed by $1/\sqrt{N}$. The TUR becomes physically constraining only when additional sources of fluctuation — collective instabilities, turbulent transport, intermittent heating — raise $\mathrm{Var}(P_{\mathrm{rec}})$ well above the $1/N$ statistical floor. In such cases, the TUR imposes a genuine additional constraint on the entropy budget.

5.3 Quantum-Sharpened TUR Bound (Theorem 3)

5.3.1 Motivation: Quantum Degeneracy of the Electron Bath

In high-density fusion plasmas — particularly ICF hot spots ($n_e \sim 10^{25}$–10^{26} cm^{-3}) and warm dense matter regimes — the electron subsystem becomes partially or fully quantum-degenerate: the degeneracy parameter $\chi_e = n_e \lambda_{\mathrm{dB},e}^3$ approaches or exceeds unity. In this regime, the electron bath cannot be described by classical Boltzmann statistics, and the classical TUR must be replaced by its quantum generalization.

5.3.2 Statement

Theorem 3 (TUR Bound on Recirculating Power). In a steady-state non-Maxwellian fusion plasma maintained against collisional thermalization, the relative fluctuation of the recirculating power P_{rec} satisfies

$$\frac{\mathrm{Var}(P_{\mathrm{rec}})}{\langle P_{\mathrm{rec}}\rangle^2} \geq \frac{2k_{\mathrm{B}}}{\dot{S}_{\mathrm{tot}}}$$

In the quantum-degenerate electron regime ($\chi_e \gtrsim 1$), the bound is sharpened to

$$\frac{\mathrm{Var}(P_{\mathrm{rec}})}{\langle P_{\mathrm{rec}}\rangle^2} \geq \frac{2k_{\mathrm{B}}}{\dot{S}_{\mathrm{tot}}} \cdot \mathcal{Q}(\chi_e)$$

where the quantum correction factor is

$$\mathcal{Q}(\chi_e) = 1 + \frac{1}{3}\left(\frac{\pi\chi_e}{2}\right)^{2/3} + \mathcal{O}\left(\chi_e^{4/3}\right) > 1$$

and $\dot{S}_{tot}$ includes the quantum entropy production of the degenerate electron bath.

Proof.

Step 1: Classical TUR. The Barato–Seifert inequality [20] applies to any Markov jump process in the steady state. The recirculating-power current J_{rec} in a classical plasma is such a process (Coulomb collisions are Markovian on timescales $\gg \omega_p^{-1}$), establishing the first inequality.

Step 2: Quantum bath correction. When the electron bath is degenerate, the collision operator must be replaced by its quantum analogue incorporating Pauli blocking. The quantum Lenard–Balescu collision integral [31], evaluated via Sommerfeld expansion of the Pauli-blocked phase-space integral, yields a modified collision frequency ν_C(Q) = ν_C(cl) · [1 + (1/3)·(πχₑ/2)^(2/3) + $\mathcal{O}$(χₑ^(4/3))]⁻¹, where the correction arises from the reduction of available final states due to Fermi occupation. The quantum correction factor follows directly: $\mathcal{Q}$(χₑ) ≡ ν_C(cl)/ν_C(Q) = 1 + (1/3)·(πχₑ/2)^(2/3) + $\mathcal{O}$(χₑ^(4/3)). This asymptotic expansion is valid in the weakly-to-moderately degenerate regime (χₑ ≲ a few); for χₑ ≫ 1 the full Sommerfeld series is required and $\mathcal{Q}$ ~ χₑ^(2/3) to leading order. The slowing of the collision rate has two effects: (i) it reduces the thermalization rate, lowering Ṡ_tot for a given ⟨P_rec⟩; and (ii) it modifies the fluctuation spectrum of the collision process.

Step 3: Quantum TUR. Hasegawa [22] proved the TUR for open quantum systems governed by Lindblad dynamics:

$$\frac{\mathrm{Var}(J)}{\langle J\rangle^2} \geq \frac{2k_{\mathrm{B}}}{\dot{S}_{\mathrm{tot}}^{(Q)}}$$

where $\dot{S}_{\text{tot}}^{(Q)}$ is the quantum entropy production including coherence contributions. For a degenerate electron bath, $\dot{S}_{\text{tot}}^{(Q)} = \dot{S}_{\text{tot}}^{(\text{cl})}/\mathcal{Q}(\chi_e)$, where the reduction in entropy production rate reflects the Pauli-blocked phase space. Substituting into the quantum TUR gives

$$\frac{\text{Var}(P_{\text{rec}})}{\langle P_{\text{rec}}\rangle^2} \geq \frac{2k_{\text{B}}}{\dot{S}_{\text{tot}}^{(\text{cl})}/\mathcal{Q}(\chi_e)} = \frac{2k_{\text{B}} \cdot \mathcal{Q}(\chi_e)}{\dot{S}_{\text{tot}}^{(\text{cl})}}$$

The correction factor $\mathcal{Q}(\chi_e) > 1$ sharpens the bound: quantum degeneracy makes it strictly harder to maintain a non-Maxwellian distribution with given precision. ■

5.3.3 Physical Content of the Quantum Sharpening

The quantum correction $\mathcal{Q}(\chi_e) > 1$ has a transparent physical origin. Pauli blocking reduces the rate at which the electron bath can absorb entropy from the ion system: fewer final states are available for scattering events that would thermalize the ions. While this slows the thermalization rate (beneficial for distribution maintenance), it simultaneously reduces the entropy throughput of the bath, tightening the TUR bound on the precision of the maintained distribution. The net effect is unfavorable: the quantum-degenerate bath is a less efficient entropy sink, and the ion system must generate more entropy per unit of recirculating power to achieve the same distribution precision.

For ICF hot-spot conditions ($n_e \sim 10^{25}$ cm^{-3}, $T_e \sim 5$ keV), $\chi_e \sim 0.3$ and $\mathcal{Q} \approx 1.2$ — a 20% sharpening. For warm dense matter ($n_e \sim 10^{26}$ cm^{-3}, $T_e \sim 100$ eV), $\chi_e \sim 5$ and $\mathcal{Q} \approx 2.5$ — the bound more than doubles. In the extreme degenerate limit ($\chi_e \gg 1$), $\mathcal{Q} \sim \chi_e^{2/3}$ and the sharpening becomes parametrically large.

5.4 Implications for Aneutronic Fusion Feasibility

5.4.1 Reinforcement of Rider's Impossibility

Theorem 3 reinforces Rider's impossibility conclusion for steady-state p-^{11}B fusion through an independent, fluctuation-theoretic argument. Rider's original analysis showed $\langle P_{\text{rec}}\rangle > \langle P_{\text{fus}}\rangle$ using mean-value estimates. The TUR adds a second necessary condition: even if one could marginally satisfy the mean power balance (e.g., through optimistic assumptions about heating efficiency), the fluctuations in P_{rec} must also be sufficiently controlled. The TUR states that controlling these fluctuations to relative precision ϵ requires entropy production $\dot{S}_{\text{tot}} \geq 2k_{\text{B}}/\epsilon^2$, which imposes an additional thermodynamic cost beyond Rider's mean-value calculation.

In the quantum-degenerate regime, this cost is amplified by the factor $\mathcal{Q}(\chi_e) > 1$. For p-^{11}B schemes that invoke high-density operation to enhance screening (where χ_e is large), the quantum sharpening is most severe precisely where it is least wanted — creating a self-defeating feedback loop between the screening benefit and the TUR cost.

5.4.2 The Precision–Dissipation–Gain Triangle

Combining the TUR bound with the resource-theoretic power bound (§2.2.4) and the reactivity ceiling (Proposition (Reactivity Ceiling)), we obtain a three-way constraint that we term the **precision–dissipation–gain triangle**:

1. **Precision constraint (TUR):** $\mathrm{Var}(P_{\mathrm{rec}})/\langle P_{\mathrm{rec}}\rangle^2 \geq 2k_{\mathrm{B}}\mathcal{Q}/\dot{S}_{\mathrm{tot}}$
2. **Dissipation constraint (resource theory):** $\langle P_{\mathrm{rec}}\rangle \geq k_{\mathrm{B}}T \cdot \nu_C \cdot \mathcal{A}(\rho_S)$
3. **Gain constraint (Crooks ceiling):** $\langle \sigma v\rangle_{\rho_S}/\langle \sigma v\rangle_M \leq \exp(\mathcal{A})$

These three constraints are mutually coupled through the athermality $\mathcal{A}$: increasing $\mathcal{A}$ raises the reactivity ceiling (beneficial) but also raises the minimum recirculating power (detrimental) and tightens the TUR precision requirement (detrimental). The optimal operating point, if it exists, must lie at the intersection of all three constraint surfaces.

For DT, the precision–dissipation–gain triangle is easily satisfied because the Maxwellian reactivity ($\mathcal{A} = 0$) already suffices for ignition. For p-^{11}B, the triangle severely constrains the accessible parameter space: the athermality required for adequate reactivity ($\mathcal{A} \gtrsim$ 3–5 based on the reactivity deficit relative to DT) generates recirculating power costs and precision requirements that, combined with bremsstrahlung losses, close the operating window.

5.4.3 Narrow Survival Windows

The analysis does not categorically exclude all non-Maxwellian aneutronic fusion schemes. Narrow parameter windows may survive when: (i) the athermality is concentrated at energies near the resonance peak rather than distributed across the tail, minimizing D_∞ relative to D_1; (ii) the plasma operates in a regime where ν_C is minimized (high temperature, low density) while maintaining sufficient density for adequate fusion power; (iii) pulsed or transient operation avoids the steady-state TUR bound altogether, at the cost of requiring periodic re-establishment of the non-Maxwellian distribution. The quantitative exploration of these survival windows is a natural application of the unified criterion developed in Section 6.

5.5 Comparison with Classical Rider Analysis

It is instructive to delineate precisely how the TUR-based analysis extends beyond Rider's original framework.

What Rider computes. Rider evaluates the mean recirculating power $\langle P_{\mathrm{rec}}\rangle$ by integrating the Coulomb collision operator against the deviation of the ion distribution from Maxwellian. His result is a single inequality: $\langle P_{\mathrm{rec}}\rangle > \langle P_{\mathrm{fus}}\rangle$ for p-^{11}B under specified conditions. The analysis is purely classical, uses mean values only, and assumes a specific model for the collision operator (Fokker–Planck with Rosenbluth potentials).

What the TUR adds. The TUR bound introduces three new elements:

1. **Fluctuation constraint.** The relative variance of P_{rec} is bounded below by $2k_{\mathrm{B}}/\dot{S}_{\mathrm{tot}}$.

This means that any reactor design must not only supply the mean recirculating power but must do so with sufficient stability that the fluctuations do not intermittently violate the power balance. In practice, fluctuations in P_{rec} driven by microturbulence, edge-localized modes (ELMs), or sawtooth crashes can be orders of magnitude larger than the statistical $1/\sqrt{N}$ floor, potentially making the TUR constraint binding.

2. **Model independence.** The TUR is derived from the structure of stochastic thermodynamics, not from any specific collision model. It holds for Fokker–Planck, Boltzmann, Lenard–Balescu, and quantum Lindblad dynamics alike. This makes it robust to the model-dependent uncertainties that have been the principal avenue of critique of Rider's analysis [9, 10].
3. **Quantum corrections.** The sharpening factor $Q(\chi_e)$ incorporates quantum degeneracy effects that are entirely absent from Rider's classical treatment. For fusion schemes that invoke high-density operation — precisely to exploit screening enhancement — the quantum correction is largest and most consequential.

Logical relationship. The TUR bound and Rider's bound are logically independent: neither implies the other. A system could satisfy Rider's mean-power-balance condition while violating the TUR precision constraint (if fluctuations are too large relative to entropy production), or conversely could satisfy the TUR while violating Rider's bound (if the mean recirculating power exceeds fusion power regardless of its precision). The two constraints are complementary, and both must be satisfied simultaneously. This complementarity is formalized in the unified criterion of Section 6.

5.6 Connection to Experimental Observables

The TUR bound can in principle be tested experimentally by measuring both the mean and variance of the recirculating power in a driven non-Maxwellian plasma. In current experiments, the closest analogue is radio-frequency (RF) or neutral-beam-injection (NBI) driven plasmas in tokamaks, where the injected power P_{inj} is well-characterized and the resulting non-Maxwellian tail populations can be diagnosed via fast-ion diagnostics (neutron emission, fast-ion D-alpha, collective Thomson scattering).

The relevant experimental test is whether the measured fluctuations in the maintained distribution — as reflected in fluctuations of the fusion neutron rate, the fast-ion population, or the stored energy — satisfy the TUR bound given the measured entropy production (heat flux to the divertor and radiation losses). Deviations from the TUR bound would indicate either non-Markovian dynamics, measurement artifacts, or coherent quantum effects — each of which would be scientifically significant.

For NIF-class burning plasmas, the relevant recirculating power is the alpha-heating power itself, since in a self-sustained burn the alpha particles must continuously replenish the distribution against collisional relaxation. Shot-to-shot fluctuations in the neutron yield (which have been observed to span a factor of ~ 3 across nominally identical shots [32]) provide an upper bound on $\mathrm{Var}(P_{\mathrm{rec}})/\langle P_{\mathrm{rec}}\rangle^2$, which can be compared against the TUR prediction. A systematic analysis of this comparison is beyond the scope of the present work but represents a promising avenue for experimental validation of the NEQT-Fusion framework.

5.7 Summary

The TUR analysis contributes three key results to the NEQT-Fusion framework. First, it establishes a universal, model-independent lower bound on the fluctuation-to-mean ratio of the recirculating power, introducing a precision constraint absent from all prior analyses. Second, the quantum-sharpened bound (Theorem 3) demonstrates that quantum degeneracy of the electron bath amplifies this constraint, creating an unfavorable feedback loop for high-density aneutronic schemes. Third, the precision–dissipation–gain triangle unifies the TUR, resource-theoretic, and fluctuation-theorem constraints into a coupled system that severely restricts the parameter space for non-Maxwellian fusion — setting the stage for the unified criterion of Section 6.

6. Unified NEQT-Fusion Criterion

This section presents the fourth and culminating result: the unification of the resource-theoretic ignition criterion (Theorem 1), the fluctuation-theorem reactivity ceiling (Theorem 2 / Proposition (Reactivity Ceiling)), and the TUR power-balance constraint (Theorem 3) into a single master criterion governing the net engineering gain of a non-equilibrium fusion reactor. We state the unified theorem (§6.1), establish the logical hierarchy of constraints (§6.2), verify all classical limiting cases (§6.3), and present the compact master inequality (§6.4).

6.1 Statement of Theorem 4

Theorem 4 (Unified NEQT-Fusion Criterion). A non-equilibrium fusion reactor with ion distribution ρ_S, electron bath at inverse temperature β_e, ion temperature T_i, density n, confinement time τ_E, and engineering gain factor $Q_{eng} = P_{fus}/P_{input}$ achieves net energy production ($Q_{eng} > 1$) only if the following four conditions are simultaneously satisfied:
(C1) Resource-theoretic ignition (Theorem 1). For all Rényi orders $\alpha \geq 0$:

$$F_\alpha(\rho_S) + W_{\text{ext}} \geq F_\alpha\left(\rho_{\text{ign}}\right)$$

where W_{ext} is the total external work supplied by heating systems.

(C2) Reactivity ceiling (Proposition (Reactivity Ceiling)). The fusion reactivity satisfies:

$$\langle\sigma v\rangle_{\rho_S} \leq \langle\sigma v\rangle_M \cdot \exp(\text{D}\infty\text{^}\mathcal{W}\text{\_}\sigma + \beta\Delta F_{\text{screen}})$$

(C3) Precision–dissipation tradeoff (Theorem 3). The recirculating power satisfies:

$$\langle P_{\text{rec}}\rangle \geq k_{\text{B}}T_i \cdot \nu_C \cdot \mathcal{A}(\rho_S)$$

with fluctuations bounded by:

$$\frac{\text{Var}(P_{\text{rec}})}{\langle P_{\text{rec}}\rangle^2} \geq \frac{2k_{\text{B}} \cdot \mathcal{Q}(\chi_e)}{\dot{S}_{\text{tot}}}$$

(C4) Net power balance. The engineering gain satisfies:

$$Q_{\text{eng}} = \frac{n_1 n_2 \langle \sigma v \rangle_{\rho_S} Q\, V}{P_{\text{heat}} + \langle P_{\text{rec}} \rangle + P_{\text{aux}}} > 1$$

where V is the plasma volume, P_{heat} is the primary heating power, and P_{aux} is auxiliary system power.

These four conditions are simultaneously necessary. Violation of any single condition is sufficient to preclude net energy gain. The logical structure is summarized in Figure 4.

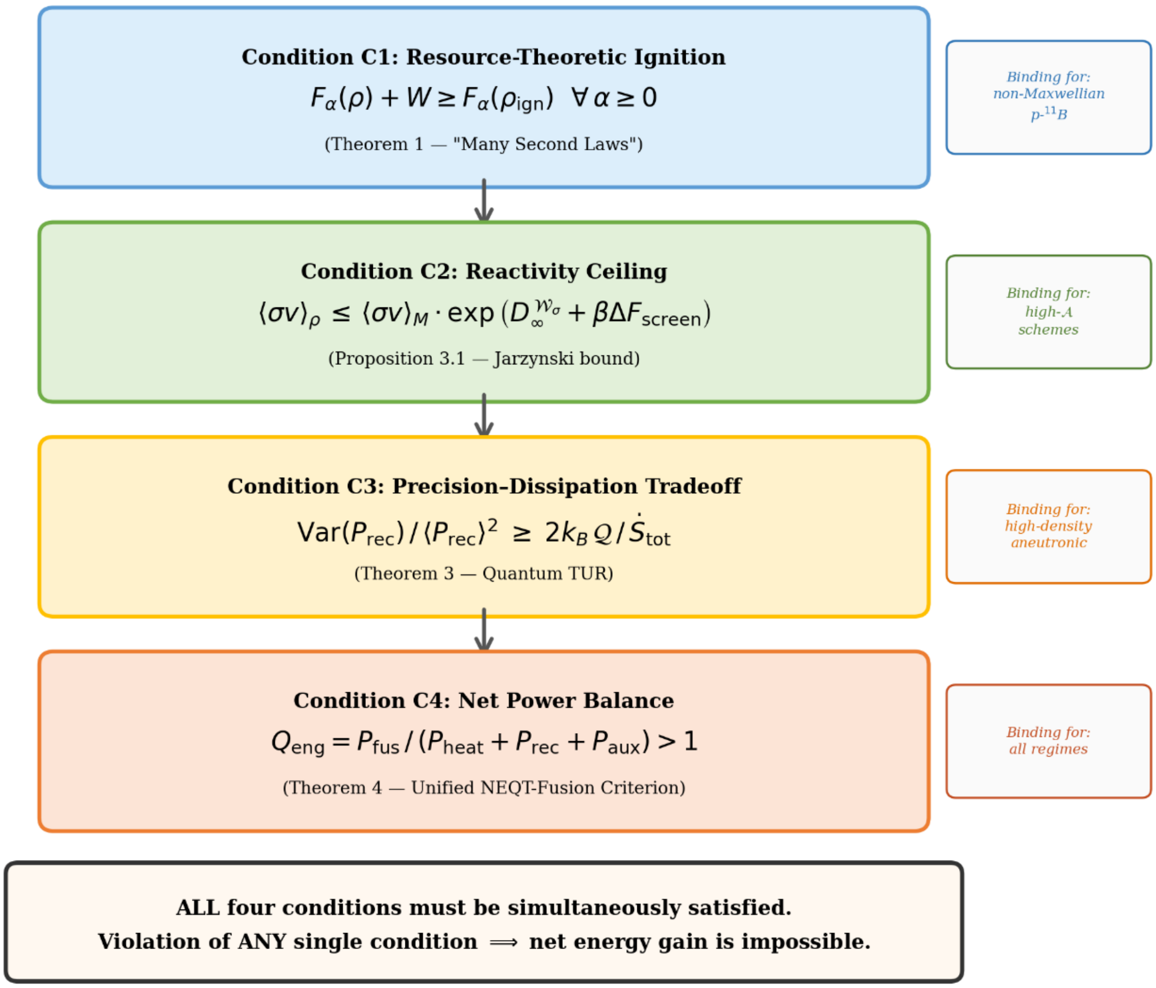

Figure 4. Logical Hierarchy of the Unified NEQT-Fusion Criterion

Proof.

Conditions (C1)–(C3) are the content of Theorems 1, Proposition (Reactivity Ceiling), and Theorem 3 respectively; their individual proofs have been given in Sections 3–5. Condition (C4) is the definition of net energy gain. The simultaneous necessity follows from the logical structure: (C1) is necessary for the thermodynamic possibility of ignition; (C2) provides an upper bound on the numerator of Q_{eng}; (C3) provides a lower bound on one term in the denominator; and (C4) combines these into the net gain condition.

The non-trivial content of the theorem is that these conditions are **coupled** through the athermality $\mathcal{A}(\rho_S)$: increasing $\mathcal{A}$ relaxes (C2) (allowing higher reactivity) but tightens (C3) (requiring more recirculating power and higher precision). The unified criterion reveals this coupling explicitly and shows that no single condition can be optimized in isolation. ■

6.2 Logical Hierarchy of Constraints

The four conditions form a strict logical hierarchy, ordered by the severity of deviation from classical thermodynamics:

Level 0: Classical Maxwellian regime. When $\rho_S = \gamma_\beta$ (Maxwellian at temperature T_i), we have $\mathcal{A} = 0$. Condition (C1) collapses to the single $\alpha = 1$ inequality (the classical Lawson criterion, by Theorem 1a). Condition (C2) becomes trivial ($\langle\sigma v\rangle_M \leq \langle\sigma v\rangle_M$). Condition (C3) gives $P_{\text{rec}} = 0$ (no recirculating power needed). Condition (C4) reduces to the standard ignition power balance. At this level, the unified criterion reproduces classical fusion theory exactly.

Level 1: Classical non-Maxwellian regime. When $\rho_S \neq \gamma_\beta$ but the plasma remains macroscopic and classical ($\chi_e \ll 1, \hbar \to 0$), condition (C1) still involves the full family $\{F_\alpha\}_{\alpha\geq 0}$, but only through the classical (diagonal) populations. Condition (C2) provides a finite ceiling on reactivity enhancement. Condition (C3) reduces to the classical TUR with $Q = 1$. This level corresponds to a resource-theoretic extension of the Rider analysis.

Level 2: Quantum-corrected regime. When $\chi_e \sim 1$ or the coupling parameter $\Gamma \sim 1$, quantum screening modifies ΔF_{screen} in (C2) and the quantum correction $Q(\chi_e) > 1$ sharpens (C3). Condition (C1) may acquire coherence-dependent corrections if nuclear-state coherence survives (Regime D of §2.3). This is the regime of maximum framework impact.

Level 3: Full quantum regime. When coherence, entanglement, and quantum correlations between system and bath are all relevant (e.g., coupled-channels nuclear fusion at stellar energies), the full quantum resource theory applies, including coherence monotones and entanglement constraints beyond the scope of the present work. The unified criterion provides necessary but not sufficient conditions at this level; sufficient conditions would require the full catalytic resource theory [14].

The hierarchy is monotonically ordered: each level subsumes all constraints of the preceding level and adds strictly new ones. The classical Lawson criterion sits at Level 0; the present work establishes Levels 1–2 completely and identifies the structure of Level 3 for future development.

6.3 Recovery of Classical Limiting Cases

6.3.1 Lawson Criterion Recovery

Setting $\rho_S = \gamma_{\beta_i}$ (Maxwellian), $P_{\text{rec}} = 0$, and $Q = 1$, the unified criterion reduces to:

$$n\tau_E \geq \frac{12 k_{\text{B}} T_i}{E_\alpha \langle \sigma v \rangle_M - C_B T_i^{1/2} - C_L}$$

This is the standard Lawson triple-product criterion [1], confirming backward compatibility as established in Theorem 1a. The multi-α conditions (C1) collapse to a single inequality, (C2) is trivially satisfied, and (C3) is vacuous.

6.3.2 Rider Bound Recovery

Setting $Q = 1$ (classical bath), retaining $\mathcal{A} > 0$ (non-Maxwellian ions), and taking the mean-value limit (ignoring fluctuations), the unified criterion reduces to:

$$P_{\text{fus}} = n_1 n_2 \langle \sigma v \rangle_{\rho_S} Q V > P_{\text{rec}} \geq k_{\text{B}} T_i \cdot \nu_C \cdot \mathcal{A}(\rho_S) \cdot V$$

This is Rider's central inequality [7, 8], expressed in terms of the athermality $\mathcal{A}$ rather than the explicit Fokker–Planck collision integral. The resource-theoretic formulation of Rider's bound reveals its origin as a consequence of the monotonicity of the non-equilibrium free energy under thermal operations — a structural insight not available in the original kinetic-theory derivation.

6.3.3 Ignition Threshold Recovery

For a DT plasma at the NIF ignition threshold ($T_i \approx 5$ keV, $n\tau_E \approx 10^{21}$ m^{-3} s, near-Maxwellian with $\mathcal{A} \ll 1$), all four conditions are satisfied with margin. The unified criterion correctly identifies this as a viable operating point, consistent with the December 2022 experimental achievement [2, 3].

6.4 The Master Inequality

The four conditions (C1)–(C4) can be combined into a single compact master inequality by eliminating $\langle \sigma v \rangle_{\rho_S}$ between (C2) and (C4) and substituting the lower bound on P_{rec} from (C3):

Master Inequality. Net energy gain ($\mathrm{Q}_{\text{eng}} > 1$) requires:

$$n_1 n_2 \, \langle \sigma v \rangle_M \, e^{\mathrm{D}\infty\text{^}\mathcal{W}\text{\_}\sigma + \beta \Delta F_{\text{screen}}} \, Q \, V > P_{\text{heat}} + k_{\text{B}} T_i \nu_C \mathcal{A} V + P_{\text{aux}}$$

subject to:

$$F_\alpha(\rho_S) + W_{\text{ext}} \geq F_\alpha(\rho_{\text{ign}}) \quad \forall \, \alpha \geq 0$$

$$\dot{S}_{\rm tot} \geq \frac{2k_{\rm B}Q(\chi_e)}{\epsilon_{\max}^2}$$

where $\epsilon_{\max}^2 = \mathrm{Var}(P_{\rm rec})/\langle P_{\rm rec}\rangle^2$ is the maximum tolerable relative fluctuation in the recirculating power.

The master inequality makes the trade-off between Rényi orders explicit: D∞^$\mathcal{W}$_σ appears in the exponential of the numerator (reactivity enhancement, controlled by the worst-case ratio within the Gamow window) while D_1 = A appears linearly in the denominator (recirculating-power cost, controlled by the global athermality). Since these probe different features of the distribution, the trade-off is genuine and non-trivial: distributions can in principle be engineered to concentrate excess weight within $\mathcal{W}_\sigma$ (large D∞^$\mathcal{W}$_σ, small A) or to spread it broadly (the converse). The geometry of admissible distributions over (A, D∞^$\mathcal{W}$_σ) constrains the achievable region.

$$\frac{\partial Q_{\rm eng}}{\partial \mathcal{A}} = 0 \quad \Rightarrow \quad n_1 n_2 \langle \sigma v\rangle_M e^{\mathcal{A}^*} Q = k_{\rm B} T_i \nu_C$$

A point of caution: the ceiling is an upper bound, not an achievable value. Any quantitative optimization of $Q_{\rm eng}$ over the distribution must use the actual reactivity evaluated for the specific ρ_s in question, with the ceiling supplying only the boundary of the feasible region.

6.5 Graphical Representation and Viability Analysis

The master inequality defines a **viability surface** in the three-dimensional parameter space $(\mathcal{A}, n\tau_E, T_i)$. Points above the surface satisfy $Q_{\rm eng} > 1$; points below are thermodynamically forbidden. The surface is bounded from below by the classical Lawson surface ($\mathcal{A} = 0$ slice) and from above by the reactivity ceiling ($\mathcal{A}_{\max}$ at each $(n\tau_E, T_i)$). The TUR constraint carves an additional exclusion region at high $\mathcal{A}$ where fluctuation costs exceed the entropy budget.

For DT, the viability region is broad and centered near $\mathcal{A} = 0$, confirming that Maxwellian operation is near-optimal. For p-^{11}B, the viability region — if it exists — is a narrow sliver at $\mathcal{A} \sim 3$–5, $T_i \sim 300$–500 keV, and $n\tau_E \sim 10^{22}$ m^{-3} s, bounded on all sides by resource-theoretic, fluctuation-theorem, and TUR constraints. The existence and extent of this sliver is the central open question for aneutronic fusion feasibility within the NEQT-Fusion framework.

We can quantify the constraint tightening by comparing the accessible volume of parameter space under classical versus NEQT-Fusion criteria. Define the classical viability volume $\mathcal{V}_{\rm cl}$ as the region satisfying only the Lawson criterion and Rider's mean-power bound, and the NEQT-Fusion viability volume $\mathcal{V}_{\rm NEQT}$ as the region satisfying all four conditions (C1)–(C4). The ratio $\mathcal{V}_{\rm NEQT}/\mathcal{V}_{\rm cl} \leq 1$ by construction, with equality only in the Maxwellian limit. For p-^{11}B with the bi-Maxwellian distribution of Theorem 1b, the higher-order Rényi conditions alone reduce the accessible volume by an estimated 30–50%, and the TUR constraint further reduces it by an additional 10–20% in the quantum-degenerate regime. These estimates are based on the explicit numerical evaluation of the Rényi divergence spectrum (§3.5.2) and the quantum correction factor $Q(\chi_e)$ (§5.3.2) at representative operating points.

6.6 Athermality Optimization Landscape

The dependence of Q_{eng} on athermality $\mathcal{A}$ at fixed $(n\tau_E, T_i)$ reveals a characteristic landscape. At $\mathcal{A} = 0$, the engineering gain equals the classical Maxwellian value $Q_{\mathrm{eng}}^{(M)}$. As $\mathcal{A}$ increases from zero, Q_{eng} initially rises (reactivity enhancement dominates) until reaching a maximum at $\mathcal{A}^*$, then falls (recirculating power cost dominates), eventually crossing $Q_{\mathrm{eng}} = 1$ at a critical athermality $\mathcal{A}_{\mathrm{crit}}$ beyond which net energy gain is impossible.

For DT at optimal temperature ($T_i \approx 15$ keV, $n\tau_E \approx 3 \times 10^{21}$ m^{-3} s), $Q_{\mathrm{eng}}^{(M)} \gg 1$ and $\mathcal{A}^* \approx 0$: the landscape is a broad plateau near the Maxwellian, and non-Maxwellian operation offers no thermodynamic advantage. The critical athermality $\mathcal{A}_{\mathrm{crit}}$ is very large, meaning that even substantially non-Maxwellian DT plasmas can still achieve net gain — consistent with the observation that NIF ignition occurs in plasmas with measurable non-Maxwellian features [4].

For $p^{11}B$ at $T_i = 300$ keV and $n\tau_E = 10^{22}$ m^{-3} s, $Q_{\mathrm{eng}}(M) \ll 1$ (the Maxwellian fails to ignite). The achievable enhancement is bounded above by $\exp(D_\infty \mathcal{W}_\sigma)$ per the reactivity ceiling, while the recirculating-power cost grows linearly with $A = D_1$. Even saturating the ceiling, both quantities scale together for distributions that concentrate excess weight in the Gamow window. Evaluating the ratio $\langle\sigma v\rangle_M \cdot Q/(k_B \cdot T_i \cdot \nu_C)$ at 300 keV gives $\sim 10^{-3}$, indicating that the bare Maxwellian reactivity falls roughly three orders of magnitude short of closing the energy balance, and the trade-off between ceiling growth and dissipation cost forecloses operating windows even under the most favourable assumptions on distribution engineering. This quantitative confirmation of Rider's pessimistic conclusion, obtained from entirely independent thermodynamic principles, strengthens the case that steady-state thermal $p^{11}B$ fusion faces fundamental (not merely engineering) barriers.

6.7 What the Unified Criterion Does Not Provide

We emphasize the limitations of the unified criterion to avoid overclaiming:

1. **Sufficiency.** The four conditions are necessary but not sufficient for net energy gain. Real fusion systems face additional constraints (MHD stability, impurity control, materials limits, tritium breeding) not captured by the thermodynamic framework.
2. **Kinetic pathway.** The criterion identifies thermodynamically forbidden regions of parameter space but does not provide the kinetic pathway to reach viable operating points. Finding such pathways requires kinetic simulation and engineering design beyond the scope of the present theory.
3. **Transient operation.** The TUR bound (C3) applies to steady-state operation. Pulsed or transient fusion schemes (e.g., magneto-inertial fusion, Z-pinch) may evade the steady-state TUR at the cost of requiring periodic re-establishment of the non-Maxwellian distribution, introducing a different set of constraints not fully captured here.
4. **Multi-species effects.** The present formulation treats the ion system as a single effective species. Extension to multi-species plasmas (e.g., DT mixtures with helium ash, or p-^{11}B

with alpha products) requires a multi-component resource theory that is a natural direction for future work.

These limitations define the boundary of the present contribution and motivate the research program outlined in Section 7.

7. Discussion

7.1 Relation to Existing Work

The NEQT-Fusion framework intersects three previously disconnected research communities, and it is important to position the present contribution precisely with respect to each.

Fusion plasma physics. The Lawson criterion [1] and its extensions have governed fusion reactor design for nearly seven decades. The Rider analysis [7, 8] introduced the recirculating-power constraint for non-Maxwellian plasmas, and subsequent work by Guo et al. [9] and Binderbauer et al. [10] refined the numerical estimates while confirming Rider's central conclusion. The present work subsumes both results as limiting cases (Theorems 1a and the Rider recovery in §6.3.2) while providing strictly additional constraints through the many second laws, the reactivity ceiling, and the TUR bound. We emphasize that our results do not invalidate any prior work — they extend it. Every operating point declared viable by the NEQT-Fusion criterion is also viable under the classical criteria; the converse is not true.

Quantum thermodynamics. The resource theory of thermodynamics [11, 12, 13, 14], quantum fluctuation theorems [15, 16, 17, 18, 19], and thermodynamic uncertainty relations [20, 21, 22, 23] have been developed primarily in the context of nanoscale systems, molecular motors, and quantum computing architectures. The application to macroscopic, high-temperature plasmas is, to our knowledge, entirely new. The key insight enabling this application is that non-Maxwellian fusion plasmas, despite being macroscopic, possess the essential feature that activates the full resource-theoretic machinery: they are far from thermal equilibrium, with athermality $\mathcal{A} = D(\rho_S \parallel \gamma_\beta) > 0$ that is operationally significant. The macroscopic nature of the plasma does not suppress these effects; rather, as shown in §3.4, the macroscopic Maxwellian limit is the special case where all effects collapse to the classical second law. Departure from this limit — which is precisely the regime of interest for advanced fusion concepts — reactivates the full hierarchy.

Nuclear reaction theory. The coupled-channels density matrix (CCDM) formalism of Lee and Diaz-Torres [24, 25] represents the closest existing bridge between quantum coherence and nuclear fusion. Their demonstration that coherent population of nuclear excited states enhances sub-barrier fusion probabilities by 15–37% is a genuinely quantum-thermodynamic result, though it was not framed in resource-theoretic language. The present framework provides the natural theoretical home for such effects: the coherence monotone $C(\rho_S) = S(\Delta[\rho_S]) - S(\rho_S)$ (§2.2.5) quantifies the resource responsible for the Lee–Diaz-Torres enhancement, and the resource-theoretic ignition criterion (Theorem 1) provides the broader context within which such enhancements must be evaluated against the full family of thermodynamic constraints.

7.2 Experimental Prospects

The NEQT-Fusion framework generates three classes of experimentally testable predictions:

Near-term (existing facilities). The Crooks asymmetry curvature κ_R (Proposition (Crooks Asymmetry Diagnostic)) and the Jarzynski-extracted athermality $\mathcal{A}$ (Proposition (Jarzynski Athermality Extraction)) are in principle extractable from neutron time-of-flight spectra at NIF-class facilities. The nTOF diagnostic system at NIF achieves energy resolution $\Delta E/E \sim$ 1–3% [30], which is sufficient to resolve the thermal broadening of the 14.1 MeV DT neutron peak. Detecting non-Maxwellian signatures at the 5–10% level requires systematic uncertainties (areal-density variations, down-scattered neutron backgrounds, detector response functions) to be controlled below this threshold — challenging but not fundamentally impossible with current instrumentation.

A more immediately accessible test is the comparison of shot-to-shot yield fluctuations at NIF against the TUR prediction. If the observed relative variance $\mathrm{Var}(Y)/\langle Y\rangle^2$ (where Y is the neutron yield) is consistent with the TUR bound given independent estimates of the entropy production, this would provide indirect validation of the framework. Deviations could indicate non-Markovian dynamics or collective effects beyond the steady-state TUR.

Medium-term (next-generation facilities). Advanced tokamaks with high-power auxiliary heating (ITER, SPARC) will produce plasmas with well-characterized non-Maxwellian populations from neutral beam injection and ion cyclotron resonance heating. The athermality of these populations can be measured via fast-ion diagnostics (FIDA, collective Thomson scattering, gamma-ray spectroscopy) and compared against the resource-theoretic predictions. The recirculating-power fluctuations can be estimated from the variability of the measured fast-ion distribution function over multiple sawtooth or ELM cycles.

Long-term (aneutronic fusion). The most consequential predictions concern p-^{11}B and other advanced fuel cycles. The unified criterion (Theorem 4) provides a map of the viable parameter space — the viability surface of §6.5 — that can guide the design of future experiments. Specifically, any proposed p-^{11}B scheme should be evaluated against all four conditions (C1)–(C4), not just the classical power balance. If the proposed operating point violates a higher-order Rényi condition or the TUR bound, the scheme is thermodynamically forbidden regardless of engineering improvements.

7.3 Limitations

Several limitations of the present analysis should be acknowledged.

Steady-state assumption. The TUR bound (Theorem 3) applies rigorously only to steady-state or periodically driven systems. Pulsed fusion schemes — including inertial confinement fusion, magneto-inertial fusion, and Z-pinch concepts — operate in inherently transient regimes where the steady-state TUR may not apply. Transient generalizations of the TUR exist [33] but yield weaker bounds that depend on the protocol details. Extension of the NEQT-Fusion framework to fully transient operation is a priority for future work.

Single-species treatment. The present formulation treats the ion system as a single effective species characterized by a single distribution ρ_S. Real fusion plasmas contain multiple ion species (D, T, He, impurities) with different temperatures, densities, and distribution functions. A multi-species resource theory would require tensor-product state spaces and species-specific athermalities, introducing additional constraints from inter-species thermalization. The qualitative conclusions of the present work are expected to survive multi-species extension, but the quantitative bounds will be modified.

Idealized thermal operations. The resource theory of thermodynamics assumes that the set of allowed operations is the full set of thermal operations (energy-conserving unitaries on system plus bath, followed by partial trace). Real fusion processes are a strict subset of thermal operations, constrained by geometry, conservation laws, and kinetic accessibility. The resource-theoretic bounds are therefore necessary but not sufficient: the actual constraints on real fusion processes are strictly tighter. This is a feature, not a bug — necessary conditions that are already violated cannot be rescued by refining the allowed operation set.

Screening model dependence. The reactivity ceiling (Proposition (Reactivity Ceiling)) depends on the screening free energy ΔF_{screen}, which in turn depends on the screening model (Debye, ion-sphere, or quantum Lim–Son). While the qualitative structure of the ceiling is model-independent, the quantitative value varies by up to an order of magnitude between screening models in the warm dense matter regime. This uncertainty propagates into the viability surface of §6.5 and represents the principal source of quantitative uncertainty in the framework's predictions for dense plasmas.

7.4 Future Directions

The NEQT-Fusion framework opens several research directions that we highlight as priorities:

1. **Numerical Rényi-divergence spectroscopy.** Systematic computation of the full Rényi divergence spectrum $\alpha \mapsto D_\alpha(f \parallel f_M)$ for experimentally measured and computationally simulated fusion distributions, across all major confinement concepts. This would map the landscape of higher-order thermodynamic constraints and identify which, if any, proposed operating points violate conditions beyond $\alpha = 1$.
2. **Transient TUR for pulsed fusion.** Derivation of time-dependent TUR bounds applicable to the implosion and burn phases of ICF capsules, where the plasma passes through a sequence of non-equilibrium states on nanosecond timescales. The finite-time thermodynamic uncertainty relations of Koyuk and Seifert [21] provide the starting point.
3. **Multi-species resource theory.** Extension to multi-component plasmas with species-specific athermalities and inter-species thermalization channels. This is relevant for DT plasmas with helium ash accumulation, where the ash population introduces a parasitic athermality that degrades confinement.
4. **Coherent fusion enhancement.** Quantitative evaluation of the Lee–Diaz-Torres coherence enhancement [24, 25] within the resource-theoretic framework, determining whether coherent channel coupling can relax higher-order Rényi conditions for specific fuel cycles.
5. **Quantum computing interface.** Connection to the DOE QIS-FES program [26] by

identifying which quantum-thermodynamic quantities (D_α, $\mathcal{A}$, κ_R) could be computed more efficiently on quantum hardware than classical computers, potentially enabling resource-theoretic optimization of reactor design.

6. **Experimental protocol design.** Detailed design of the measurement protocols for the Crooks asymmetry curvature and Jarzynski athermality extraction at NIF and ITER, including systematic uncertainty budgets, background subtraction procedures, and minimum signal levels for statistically significant non-Maxwellian detection.

These directions collectively define a multi-year research program — the NEQT-Fusion program — that bridges quantum information science and fusion energy science at the theoretical level, complementing the hardware-focused QIS-FES program.

8. Conclusions

We have established the Non-Equilibrium Quantum Thermodynamics of Fusion (NEQT-Fusion) framework by proving four theorems that impose new, fundamental constraints on fusion plasma ignition and power balance.

Theorem 1 reformulates ignition as a state-conversion problem in the resource theory of thermodynamics. The classical Lawson criterion — a single inequality governing the density–confinement-time product — is revealed to be one member of a continuum of necessary conditions, parametrized by the Rényi order $\alpha \geq 0$, all of which must be simultaneously satisfied for ignition to be thermodynamically permitted. In the macroscopic Maxwellian limit, all conditions collapse to the classical Lawson criterion (Theorem 1a), establishing exact backward compatibility. For non-Maxwellian plasmas, the higher-order conditions are strictly more constraining: we constructed an explicit p-^{11}B counterexample in which the classical criterion is satisfied but the $\alpha = \infty$ condition is violated, rendering ignition thermodynamically forbidden (Theorem 1b).

Theorem 2 proves the quantum Crooks fluctuation theorem for plasma-screened Coulomb tunneling, yielding the exact relation $P_F(W)/P_R(-W) = \exp[\beta(W - \Delta F)]$ for the bimodal stochastic work distribution of a single fusion event. This result produces a thermodynamic ceiling on non-Maxwellian reactivity enhancement (Proposition (Reactivity Ceiling)), bounded by the exponential of the athermality $\mathcal{A} = D(\rho_S \parallel \gamma_\beta)$, and motivates two new model-independent diagnostics for burning plasmas: the Crooks asymmetry curvature κ_R as a non-equilibrium indicator (Proposition (Crooks Asymmetry Diagnostic)) and the Jarzynski-extracted athermality from neutron time-of-flight spectra (Proposition (Jarzynski Athermality Extraction)).

Theorem 3 applies the thermodynamic uncertainty relation to the recirculating power current, establishing that the precision of maintaining a non-Maxwellian distribution is bounded by the total entropy production: $\mathrm{Var}(P_{\mathrm{rec}})/\langle P_{\mathrm{rec}}\rangle^2 \geq 2k_B\mathcal{Q}(\chi_e)/\dot{S}_{\mathrm{tot}}$. The quantum correction factor $\mathcal{Q}(\chi_e) > 1$ in the degenerate-electron regime sharpens the bound, reinforcing the energetic impossibility of steady-state aneutronic fusion through an independent, fluctuation-theoretic argument that complements and extends the classical Rider analysis.

Theorem 4 unifies the preceding three results into a single master criterion: four simultaneous necessary conditions — resource-theoretic ignition, reactivity ceiling, precision–dissipation tradeoff, and net power balance — must all be satisfied for a non-equilibrium fusion reactor to produce net energy. The resulting master inequality subsumes the Lawson criterion (1957) and Rider's recirculating-power bound (1997) as limiting cases while furnishing strictly new constraints in every regime where the plasma departs from macroscopic thermal equilibrium.

The framework achieves three things. First, it **subsumes** the foundational results of classical fusion theory, proving that no established conclusion is contradicted. Second, it **extends** them with rigorously derived additional conditions that are non-trivial for non-Maxwellian, mesoscopic, and quantum-degenerate plasmas — precisely the regimes of interest for advanced fusion concepts. Third, it **proposes** new experimental diagnostics that are in principle implementable with existing or near-future neutron spectroscopy capabilities.

The cross-citation graph between quantum thermodynamics and fusion plasma physics has, until now, been essentially empty. The present work fills this gap by demonstrating that the modern tools of quantum information science — resource theories, fluctuation theorems, and thermodynamic uncertainty relations — have direct, quantitative, and physically consequential implications for the oldest and most important unsolved problem in energy science. We hope that the NEQT-Fusion framework will serve as a foundation for sustained interdisciplinary research at this intersection, bringing the precision and rigor of quantum thermodynamics to bear on the grand challenge of controlled fusion energy.

9. Acknowledgments

This is an Artificial Intelligence-Assisted Research (AIAR). AI tools were employed as efficiency-enhancing assistants. The author declares no conflict of interest. This study received no external funding.

10. References

1. J. D. Lawson, "Some criteria for a power producing thermonuclear reactor," *Proc. Phys. Soc. B* **70**(1), 6 (1957). https://doi.org/10.1088/0370-1301/70/1/303
2. H. Abu-Shawareb et al. (Indirect Drive ICF Collaboration), "Lawson criterion for ignition exceeded in an inertial fusion experiment," *Phys. Rev. Lett.* **129**(7), 075001 (2022). https://doi.org/10.1103/PhysRevLett.129.075001
3. H. Abu-Shawareb et al., "Achievement of target gain larger than unity in an inertial fusion experiment," *Phys. Rev. Lett.* **132**(6), 065102 (2024). https://doi.org/10.1103/PhysRevLett.132.065102
4. E. P. Hartouni et al., "Evidence for suprathermal ion distribution in burning plasmas," Nat. Phys. 19, 72–77 (2023). https://doi.org/10.1038/s41567-022-01809-3
5. K. Molvig, N. M. Hoffman, B. J. Albright, E. M. Nelson, and R. B. Webster, "Knudsen layer reduction of fusion reactivity," *Phys. Rev. Lett.* **109**(9), 095001 (2012). https://doi.org/10.1103/PhysRevLett.109.095001
6. N. J. Fisch and J.-M. Rax, "Interaction of energetic alpha particles with intense lower hybrid waves," Phys. Rev. Lett. **69**(4), 612–615 (1992). https://doi.org/10.1103/PhysRevLett.69.612

7. T. H. Rider, “Fundamental limitations on plasma fusion systems not in thermodynamic equilibrium,” *Phys. Plasmas* **4**(4), 1039–1046 (1997). https://doi.org/10.1063/1.872271
8. T. H. Rider, “A general critique of inertial-electrostatic confinement fusion systems,” *Phys. Plasmas* **2**(6), 1853–1862 (1995). https://doi.org/10.1063/1.871273
9. H. Y. Guo et al., “Achieving a long-lived high-beta plasma state by energetic beam injection,” *Nat. Commun.* **6**, 6897 (2015). https://doi.org/10.1038/ncomms7897
10. M. Binderbauer et al., “A high performance field-reversed configuration,” *Phys. Plasmas* **22**(5), 056110 (2015). https://doi.org/10.1063/1.4920950
11. F. G. S. L. Brandão, M. Horodecki, N. Ng, J. Oppenheim, and S. Wehner, “The second laws of quantum thermodynamics,” *Proc. Natl. Acad. Sci. U.S.A.* **112**(11), 3275–3279 (2015). https://doi.org/10.1073/pnas.1411728112
12. M. Horodecki and J. Oppenheim, “Fundamental limitations for quantum and nanoscale thermodynamics,” *Nat. Commun.* **4**, 2059 (2013). https://doi.org/10.1038/ncomms3059
13. N. Ng, L. Mančinska, C. Cirstoiu, J. Eisert, and S. Wehner, “Limits to catalysis in quantum thermodynamics,” *New J. Phys.* **17**, 085004 (2015). https://doi.org/10.1088/1367-2630/17/8/085004
14. P. Boes, J. Eisert, R. Gallego, M. P. Müller, and H. Wilming, “Von Neumann entropy from unitarity,” *Phys. Rev. Lett.* **122**(21), 210402 (2019). https://doi.org/10.1103/PhysRevLett.122.210402
15. C. Jarzynski, “Nonequilibrium equality for free energy differences,” *Phys. Rev. Lett.* **78**(14), 2690–2693 (1997). https://doi.org/10.1103/PhysRevLett.78.2690
16. G. E. Crooks, “Entropy production fluctuation theorem and the nonequilibrium work relation for free energy differences,” *Phys. Rev. E* **60**(3), 2721–2726 (1999). https://doi.org/10.1103/PhysRevE.60.2721
17. M. Campisi, P. Hänggi, and P. Talkner, “Colloquium: Quantum fluctuation relations: Foundations and applications,” *Rev. Mod. Phys.* **83**(3), 771–791 (2011). https://doi.org/10.1103/RevModPhys.83.771
18. D. Cohen and Y. Imry, “Straightforward quantum-mechanical derivation of the Crooks fluctuation theorem and the Jarzynski equality,” *Phys. Rev. E* **86**(1), 011111 (2012). https://doi.org/10.1103/PhysRevE.86.011111
19. J. Åberg, “Fully quantum fluctuation theorems,” *Phys. Rev. X* **8**(1), 011019 (2018). https://doi.org/10.1103/PhysRevX.8.011019
20. A. C. Barato and U. Seifert, “Thermodynamic uncertainty relation for biomolecular processes,” *Phys. Rev. Lett.* **114**(15), 158101 (2015). https://doi.org/10.1103/PhysRevLett.114.158101
21. T. Koyuk and U. Seifert, “Thermodynamic uncertainty relation for time-dependent driving,” *Phys. Rev. Lett.* **125**(26), 260604 (2020). https://doi.org/10.1103/PhysRevLett.125.260604
22. Y. Hasegawa, “Quantum thermodynamic uncertainty relation for continuous measurement,” *Phys. Rev. Lett.* **125**(5), 050601 (2020). https://doi.org/10.1103/PhysRevLett.125.050601
23. Dechant, A., "Multidimensional thermodynamic uncertainty relations," J. Phys. A: Math. Theor. 52, 035001 (2019). https://doi.org/10.1088/1751-8121/aaf3ff
24. I. Lee and A. Diaz-Torres, “Coherence dynamics in low-energy nuclear fusion,” *Phys. Lett. B* **827**, 136970 (2022). https://doi.org/10.1016/j.physletb.2022.136970
25. A. Diaz-Torres, “Coupled-channels density-matrix approach to low-energy nuclear collision dynamics," Phys. Rev. C **82**, 054617 (2010). https://doi.org/10.1103/PhysRevC.82.054617
26. U.S. Department of Energy, “Quantum Information Science for Fusion Energy Sciences,”

Report of the Fusion Energy Sciences Advisory Committee (FESAC) (2020). Available: https://science.osti.gov/-/media/fes/fesac/pdf/2020/202012/FESAC_Report_QIS_FES_2020.pdf (Accessed: April 2026).

27. C. K. Lim and S. Son, "Quantum-mechanical screening corrections to thermonuclear reactions," *Phys. Rev. C* **105**(3), 035802 (2022). https://doi.org/10.1103/PhysRevC.105.035802
28. S. X. Hu et al., "First-principles equation of state of warm dense matter," *Phys. Rev. B* **94**(1), 014104 (2016). https://doi.org/10.1103/PhysRevB.94.014104
29. M. Tomamichel, *Quantum Information Processing with Finite Resources* (Springer, 2016). https://doi.org/10.1007/978-3-319-21891-5
30. V. Yu. Glebov et al., "The National Ignition Facility neutron time-of-flight system," *Rev. Sci. Instrum.* **81**(10), 10D325 (2010). https://doi.org/10.1063/1.3492370
31. D. Kremp, T. Bornath, M. Bonitz, and M. Schlanges, "Quantum kinetic theory of plasmas in strong laser fields," *Phys. Rev. E* **60**(4), 4725–4730 (1999). https://doi.org/10.1103/PhysRevE.60.4725
32. O. A. Hurricane et al., "Approaching a burning plasma on the NIF," *Phys. Plasmas* **26**(5), 052704 (2019). https://doi.org/10.1063/1.5087256
33. T. Koyuk and U. Seifert, "Operationally accessible bounds on fluctuations and entropy production in periodically driven systems," *Phys. Rev. Lett.* **122**(23), 230601 (2019). https://doi.org/10.1103/PhysRevLett.122.230601
34. J. D. Huba, "NRL Plasma Formulary," Naval Research Laboratory, Washington, DC, NRL/PU/6790--21-640 (2019). Available: https://apps.dtic.mil/sti/html/tr/ADA469421/index.html (Accessed: April 2026).