Module 9 — Quantum Thermodynamics

Quantum Foundations

hbar.university

Thermal States and Gibbs Distributions

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.1

This is the foundation node of the quantum thermodynamics module. The Gibbs state is the unique density matrix that maximizes von Neumann entropy subject to a fixed expected energy. It is the equilibrium state in quantum statistical mechanics, and it is the structural object on which all subsequent nodes — work, heat, fluctuation theorems, Landauer's principle — are built.

Formal Definition

For a quantum system with Hamiltonian H at inverse temperature $\beta = 1/(k_B T)$, the **thermal state** (or **Gibbs state**) is

$$\rho_\beta = \frac{e^{-\beta H}}{Z(\beta)}, \qquad Z(\beta) = \mathrm{tr}(e^{-\beta H}).$$

The normalization $Z(\beta)$ is the **partition function**. In the energy eigenbasis $\{|n\rangle\}$ with eigenvalues E_n, the Gibbs state is diagonal with populations

$$p_n = \frac{e^{-\beta E_n}}{Z(\beta)}.$$

The thermodynamic free energy is

$$F(\beta) = -\frac{1}{\beta} \log Z(\beta),$$

and entropy, internal energy, and other thermodynamic quantities follow by standard derivatives.

Variational Characterization

The Gibbs state is the unique density matrix that **maximizes the von Neumann entropy** subject to a constraint on the expected energy:

$$\rho_\beta = \arg \max_{\rho : \mathrm{tr}(\rho H) = U} S(\rho).$$

The Lagrange multiplier of the energy constraint is exactly β.

Equivalently, it minimizes the **free energy functional** $F(\rho) = \mathrm{tr}(\rho H) - T\,S(\rho)$. Both characterizations are quantum analogues of the Boltzmann maximum-entropy and Helmholtz minimum-free-energy principles.

Intuition

The Gibbs state is "what you get when a quantum system equilibrates with a thermal bath at temperature T." It has two structural properties that distinguish it from other density matrices:

1. **Diagonal in the energy eigenbasis.** Energy is conserved, so coherences between eigenstates of H are unstable under thermal fluctuations and decay.
2. **Boltzmann-weighted populations.** Energetically lower states are exponentially more likely. The exponential structure is the structural fingerprint of thermal equilibrium.

The Gibbs state is also the **fixed point of any Lindblad dynamics** that satisfies detailed balance with respect to H at temperature T. This connects quantum thermodynamics to the open-system formalism of Module 8.

Structural Role

The Gibbs state organizes the entire module:

- **Quantum work and heat** (9.2) are defined as energy changes during processes that take a system from one Gibbs state to another (or that drive it out of equilibrium).
- **Fluctuation theorems** (9.3) constrain the statistics of work in non-equilibrium processes.
- **Landauer's principle** (9.4) bounds the thermodynamic cost of erasing information in terms of $k_B T \log 2$.
- **Quantum heat engines** (9.5) are cyclic processes between Gibbs states at different temperatures.
- **Resource theories** (9.6) take Gibbs states as the "free" resource and quantify deviations from equilibrium as the costly resource.

Mathematical Development

High and low temperature limits

- $T \to 0$ ($\beta \to \infty$): the Gibbs state collapses onto the ground state (or ground subspace, if degenerate). Thermal fluctuations are suppressed.
- $T \to \infty$ ($\beta \to 0$): the Gibbs state becomes maximally mixed, $\rho_\beta \to I/d$. All energy levels are equally populated.

In between, the populations interpolate exponentially.

Connection to entropy

Substituting the Gibbs state into the von Neumann entropy formula:

$$S(\rho_\beta) = -\sum_n p_n \log p_n = \beta U + \log Z = \beta(U - F).$$

This is exactly the thermodynamic relation $S = (U - F)/T$, now derived from the underlying quantum-mechanical formalism rather than postulated.

Quantum partition function

For a single qubit with $H = \frac{1}{2}\omega Z$:

$$Z(\beta) = e^{-\beta\omega/2} + e^{\beta\omega/2} = 2\cosh(\beta\omega/2),$$

$$\langle E \rangle = -\frac{\partial \log Z}{\partial \beta} = -\frac{\omega}{2}\tanh(\beta\omega/2).$$

In the high-T limit $\langle E \rangle \to 0$ (qubit unbiased); in the low-T limit $\langle E \rangle \to -\omega/2$ (qubit in ground state).

Worked Example

Two-level atom in a thermal bath

Consider a two-level atom with energies 0 and $\hbar\omega$. At temperature T, the populations are

$$p_0 = \frac{1}{1 + e^{-\beta\hbar\omega}}, \qquad p_1 = \frac{e^{-\beta\hbar\omega}}{1 + e^{-\beta\hbar\omega}}.$$

This is the standard Boltzmann distribution. As T increases, p_1 approaches $1/2$; as $T \to 0$, $p_1 \to 0$.

The von Neumann entropy is the binary entropy $H_2(p_1)$, which is exactly the classical Shannon entropy of the population distribution. This is no coincidence: when a state is diagonal in a fixed basis, von Neumann entropy reduces to Shannon entropy.

Cross-Links

- **Module 5.1:** Density matrices — Gibbs states are a particularly important family.
- **Module 7.4:** Von Neumann entropy is the Lagrange-dual quantity that the Gibbs state maximizes.
- **Module 8:** Open-system dynamics with detailed balance converge to the Gibbs state.
- **Statistics:** Maximum entropy principle in classical statistics; the Gibbs state is the quantum analogue.
- **VQA:** Variational quantum eigensolvers minimize $\langle H \rangle$, which is the $T \to 0$ limit of free-energy minimization.

Structural Compression

Invariant: The Gibbs state is the maximum-entropy state at fixed expected energy.

Mechanism: $\rho_\beta = e^{-\beta H}/Z$ — the unique density matrix with this variational property.

Cross-System Echo: Boltzmann distribution in classical statistical mechanics; maximum-entropy distributions in classical information theory. Same structural identity, non-commutative generalization.

Exercises

1. Derive the Gibbs state by maximizing $S(\rho)$ subject to $\mathrm{tr}(\rho) = 1$ and $\mathrm{tr}(\rho H) = U$, using Lagrange multipliers.
2. Compute the partition function and average energy for a harmonic oscillator with $H = \hbar\omega(a^\dagger a + 1/2)$.

3. Show that the Gibbs state is the unique fixed point of any Lindblad evolution that satisfies detailed balance with respect to H at temperature T.
4. Verify that for a diagonal density matrix, von Neumann entropy reduces to Shannon entropy of the population distribution.
5. Show that the free energy $F = U - TS$ is minimized over all density matrices by the Gibbs state, and that the minimum value equals $-k_B T \log Z$.

Quantum Work and Heat

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.2

The first law of thermodynamics splits energy change into work and heat. In quantum mechanics this split is non-trivial: there is no self-adjoint operator whose expectation value gives "work done," and the very definition depends on how the system is probed. The **two-point measurement (TPM) scheme** is the standard operational resolution. It makes work a *stochastic variable* on a probability space defined by two projective energy measurements, and it connects directly to the fluctuation theorems of 9.3 and the engine cycles of 9.5.

Formal Definition

Let a system have a time-dependent Hamiltonian $H(t)$ over $t \in [0, \tau]$, with instantaneous spectral decomposition

$$H(t) = \sum_n E_n(t)\,\Pi_n(t).$$

The **two-point measurement protocol** is:

1. Prepare the system in the initial thermal state $\rho_0 = e^{-\beta H(0)}/Z_0$.
2. Measure $H(0)$ projectively. Outcome $E_n(0)$ occurs with probability $p_n^0 = e^{-\beta E_n(0)}/Z_0$.
3. Evolve the post-measurement state under the unitary $U_\tau = \mathcal{T}\exp\big(-i\int_0^\tau H(t)\,dt\big)$ generated by the driving protocol.
4. Measure $H(\tau)$ projectively. Outcome $E_m(\tau)$ occurs with conditional probability $p_{m|n} = |\langle m_\tau|U_\tau|n_0\rangle|^2$.

The **work** done by the protocol in a single run is the random variable

$$W = E_m(\tau) - E_n(0),$$

with joint probability

$$P(n,m) = p_n^0\,p_{m|n}, \qquad P(W) = \sum_{n,m} P(n,m)\,\delta\big(W - [E_m(\tau) - E_n(0)]\big).$$

For an **isolated** driven system (no bath contact during $[0,\tau]$), all energy change is work and there is no heat:

$$\langle W\rangle = \mathrm{tr}(\rho_\tau H(\tau)) - \mathrm{tr}(\rho_0 H(0)).$$

For a system in contact with a bath, heat is identified through the bath's energy change:

$$Q = -\Delta E_{\text{bath}}, \qquad \Delta E_{\text{sys}} = W + Q.$$

This is the **quantum first law** at the level of averages, and at the level of individual trajectories if one also measures the bath.

Intuition

Classically, work is "the integral of force along a path" and heat is "energy flowing across a contact." Quantum mechanically this picture breaks because:

- There is no phase-space trajectory to integrate along.
- Energy is in general not even well-defined until it is measured (if $[\rho, H] \neq 0$).
- Any attempt to continuously monitor energy changes during the protocol disturbs the state and changes the dynamics.

The TPM scheme sidesteps these problems by only asking about energy at the *endpoints*. Between the two measurements the system evolves unitarily (or via a Lindblad map if the bath is present), and the "work" is whatever the energy meter says changed. The *protocol* — the time-dependent parameters of $H(t)$ — is the analogue of classical "pushing a piston." The unitary does not commute with the projectors in general, so work is genuinely stochastic: the same protocol produces a distribution $P(W)$, not a single number.

Heat, in the same scheme, is what the *bath* recorded. If you can monitor bath energy (idealized), Q is also a two-point measurement on H_{bath}. In practice you usually infer Q from an ensemble average and the first law.

Structural Role

9.2 is the bridge between equilibrium and non-equilibrium quantum thermodynamics. It takes the Gibbs state of 9.1 as the initial condition and asks: what happens when you drive the Hamiltonian away from equilibrium?

- **9.3 Fluctuation theorems** are identities on the work distribution $P(W)$. Jarzynski's equality $\langle e^{-\beta W} \rangle = e^{-\beta \Delta F}$ is a statement about the TPM distribution and requires exactly this definition of work.
- **9.5 Quantum heat engines** are cyclic protocols whose net work output is computed by averaging W over many TPM runs.
- **9.4 Landauer's principle** is the special case where the protocol erases a bit and the minimum average heat dissipated is $k_B T \log 2$.
- **9.6 Resource theories** take Gibbs-preserving or thermal operations as free and quantify work as the cost to move between non-equilibrium states.

Without the TPM scheme, none of these downstream nodes has a well-posed starting point.

Mathematical Development

Characteristic function of work

The characteristic function $\chi(u) = \langle e^{iuW} \rangle$ admits a compact two-time-correlation form:

$$\chi(u) = \operatorname{tr}\left[e^{iuH(\tau)}\, U_\tau\, e^{-iuH(0)}\, \rho_0\, U_\tau^\dagger \right].$$

Fourier-inverting gives $P(W)$. This representation makes the fluctuation theorems of 9.3 easy to derive: they become symmetry statements about $\chi(u)$ under $u \to -u + i\beta$.

Average work from unitary dynamics

For an isolated protocol,

$$\langle W \rangle = \sum_{n,m} p_n^0\, |\langle m_\tau | U_\tau | n_0 \rangle|^2\, (E_m(\tau) - E_n(0)) = \operatorname{tr}(\rho_\tau H(\tau)) - \operatorname{tr}(\rho_0 H(0)),$$

so $\langle W \rangle$ coincides with the *unmeasured* expected energy change. Individual runs deviate from this average, which is why work fluctuations are non-trivial even in pure unitary evolution.

Quantum Jarzynski identity (preview of 9.3)

Direct calculation of $\langle e^{-\beta W} \rangle$ using $p_n^0 = e^{-\beta E_n(0)}/Z_0$:

$$\langle e^{-\beta W} \rangle = \sum_{n,m} \frac{e^{-\beta E_n(0)}}{Z_0} \, p_{m|n} \, e^{-\beta(E_m(\tau)-E_n(0))} = \frac{1}{Z_0} \sum_m e^{-\beta E_m(\tau)} = \frac{Z_\tau}{Z_0} = e^{-\beta \Delta F}.$$

This is the Jarzynski equality, obtained as a three-line consequence of the TPM definition.

First law for Lindblad dynamics

For an open system evolving under a time-dependent Lindblad generator $\mathcal{L}_t$, differentiate the expected energy:

$$\frac{d}{dt} \mathrm{tr}(\rho_t H(t)) = \underbrace{\mathrm{tr}(\rho_t \, \partial_t H(t))}_{\dot{W}} + \underbrace{\mathrm{tr}((\mathcal{L}_t \rho_t) H(t))}_{\dot{Q}}.$$

Work is associated with the explicit time dependence of H; heat is associated with the dissipator. Integrating gives the average first law. The TPM scheme reproduces this at the level of averages while also giving access to the *distribution* of W over trajectories.

Worked Example

Driven qubit under a sudden quench

Take $H(t) = \frac{\omega}{2} Z$ for $t < 0$, then suddenly quench to $H = \frac{\omega}{2}(\cos\theta \, Z + \sin\theta \, X)$ at $t = 0$, and read out the new energy immediately. Initial state is $\rho_0 = e^{-\beta H(0)}/Z_0$.

Step 1: Initial measurement. The system is diagonal in the original Z basis, so the first measurement leaves it in $|0\rangle$ or $|1\rangle$ with probabilities

$$p_0^0 = \frac{1}{1 + e^{-\beta\omega}}, \qquad p_1^0 = \frac{e^{-\beta\omega}}{1 + e^{-\beta\omega}}.$$

(Energies $-\omega/2$ and $+\omega/2$.)

Step 2: Unitary. The quench is instantaneous, so $U_\tau = I$ and the post-measurement state is just $|0\rangle$ or $|1\rangle$.

Step 3: Final measurement. The new Hamiltonian has eigenstates $|+_\theta\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$ and $|-_\theta\rangle = -\sin(\theta/2)|0\rangle + \cos(\theta/2)|1\rangle$ with energies $\pm\omega/2$.

Overlap probabilities from $|n\rangle$ to $|\pm_\theta\rangle$:

$$p_{\pm|0} = \cos^2(\theta/2), \ \sin^2(\theta/2); \qquad p_{\pm|1} = \sin^2(\theta/2), \ \cos^2(\theta/2).$$

Step 4: Work distribution. Four outcomes with values $W \in \{0, +\omega, -\omega\}$:

- $W = 0$: start $|0\rangle$ end $|+_\theta\rangle$, or start $|1\rangle$ end $|-_\theta\rangle$.
- $W = +\omega$: start $|0\rangle$ end $|-_\theta\rangle$ (energy $-\omega/2 \to +\omega/2$).
- $W = -\omega$: start $|1\rangle$ end $|+_\theta\rangle$.

Assembling:

$$P(W=+\omega)=p_0^0\sin^2(\theta/2),\quad P(W=-\omega)=p_1^0\sin^2(\theta/2),\quad P(W=0)=\cos^2(\theta/2).$$

Check Jarzynski. Compute $\langle e^{-\beta W}\rangle$:

$$\cos^2(\theta/2)+p_0^0\sin^2(\theta/2)e^{-\beta\omega}+p_1^0\sin^2(\theta/2)e^{+\beta\omega}.$$

Substituting $p_0^0=1/(1+e^{-\beta\omega})$ and simplifying yields Z_τ/Z_0, consistent with the general identity. At $\theta=0$ (no quench) $P(W=0)=1$ and all identities reduce trivially. At $\theta=\pi$ (complete flip) the distribution is maximally broad at fixed β.

Cross-Links

- **9.1 Gibbs state** — defines the initial condition for TPM protocols.
- **9.3 Fluctuation theorems** — identities on $P(W)$; Jarzynski, Crooks, Tasaki–Crooks.
- **9.5 Quantum heat engines** — work output is the TPM expectation over a cycle.
- **8.2 Lindblad equation** — separates work (from $\partial_t H$) from heat (from dissipator) at the level of averages.
- **3.x Projective measurement** — TPM is two projective measurements of $H(0)$ and $H(\tau)$.
- **Statistics** — W is a genuine random variable; $P(W)$ is a probability distribution with finite moments.

Structural Compression

Invariant: Quantum work is a stochastic variable defined by two projective energy measurements at the endpoints of a driving protocol.

Mechanism: Two-point measurement scheme. $W=E_m(\tau)-E_n(0)$ with joint probability $p_n^0 p_{m|n}$.

Cross-System Echo: Classical work is a path integral of force. Quantum work is a difference of two measurement outcomes. The first law $\Delta E=W+Q$ survives, but W and Q lose their operator identity and become distributions.

Exercises

1. Derive the characteristic function $\chi(u)=\mathrm{tr}[e^{iuH(\tau)}U_\tau e^{-iuH(0)}\rho_0 U_\tau^\dagger]$ from the TPM joint probability, and verify that $\langle W\rangle=-i\partial_u\chi|_{u=0}$ recovers the expected first-law expression.
2. Compute the full work distribution $P(W)$ for a qubit with $H(0)=\frac{\omega}{2}Z$ quenched suddenly to $H(\tau)=\frac{\omega'}{2}Z$ (pure frequency change, no rotation). Show that $P(W)$ is supported on four points and that $\langle W\rangle=(\omega'-\omega)\langle Z\rangle/2$.
3. Show that if $[H(0),H(\tau)]=0$, the TPM work distribution coincides with the classical distribution over energy-eigenstate transitions, and $P(W)$ has no quantum-coherence contributions.
4. For a harmonic oscillator with $H(t)=\hbar\omega(t)(a^\dagger a+1/2)$ under an adiabatic frequency ramp, show that the work distribution factors into contributions from each occupation number and that $\langle W\rangle=\hbar\Delta\omega\langle n\rangle$.
5. Derive the quantum Jarzynski identity $\langle e^{-\beta W}\rangle=e^{-\beta\Delta F}$ from the TPM definition in under five lines, and identify exactly where the assumption "initial state is Gibbs" enters.

6. Consider an open-system protocol where H is held fixed but the system is put in contact with a bath. Show that the TPM scheme assigns $W = 0$ on every run, while the average energy change equals the heat absorbed from the bath.

7. **Argue, structurally**, why there cannot exist a self-adjoint "work operator" $\hat{W}$ whose expectation value $\langle \hat{W} \rangle$ in the initial state agrees with the TPM average $\langle W \rangle$ *and* whose variance agrees with the TPM variance, for all protocols. Use the fact that variances of self-adjoint operators are basis-independent but TPM variances depend on the initial-energy eigenbasis.

Fluctuation Theorems (Jarzynski, Crooks)

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.3

Fluctuation theorems are exact identities on the statistics of work done by a non-equilibrium protocol, valid arbitrarily far from equilibrium. They sharpen the second law from an inequality $\langle W \rangle \geq \Delta F$ to an equality $\langle e^{-\beta W} \rangle = e^{-\beta \Delta F}$ and, in the refined Crooks form, to a detailed-balance statement relating forward and time-reversed protocols. In the quantum setting they are consequences of the TPM scheme of 9.2 together with unitarity and time-reversal symmetry.

Formal Statement

Let the system start in a thermal state $\rho_0 = e^{-\beta H(0)}/Z_0$, evolve under a driving protocol that changes $H(t)$ from $H(0)$ to $H(\tau)$, and let W be the TPM work of 9.2. Let $\Delta F = -\beta^{-1} \log(Z_\tau / Z_0)$ be the equilibrium free-energy difference.

Quantum Jarzynski equality (Tasaki 2000, Kurchan 2000):

$$\boxed{\langle e^{-\beta W} \rangle_F = e^{-\beta \Delta F}}$$

where $\langle \cdot \rangle_F$ is the average over the forward-protocol work distribution $P_F(W)$.

Quantum Crooks fluctuation theorem:

$$\boxed{\frac{P_F(W)}{P_R(-W)} = e^{\beta(W - \Delta F)}}$$

where P_R is the work distribution under the **time-reversed protocol**, run with initial state $e^{-\beta H(\tau)}/Z_\tau$ and Hamiltonian trajectory $\tilde{H}(t) = \Theta H(\tau - t) \Theta^{-1}$, with Θ the antiunitary time-reversal operator.

Jarzynski is obtained from Crooks by integrating over W. The second law $\langle W \rangle \geq \Delta F$ follows from Jarzynski by Jensen's inequality applied to the convex function $e^{-\beta x}$.

Intuition

The second law says "on average, you can't extract more work than the free-energy difference allows." Fluctuation theorems say something stronger: the probability of a rare *violation* (an individual run where $W < \Delta F$) is *exactly* exponentially suppressed, and the suppression ratio is universal. If you see an individual run where the work is $\Delta F - k_B T$, it is $e \approx 2.7$ times less likely than the time-reversed run producing work $\Delta F - k_B T + 2k_B T = \Delta F + k_B T$.

This refines the second law into a *microscopic reversibility* principle. The underlying physics is: unitary quantum dynamics is time-reversal symmetric; the only thing that breaks symmetry is the thermal initial condition, and that asymmetry is exactly Gibbs-weighted. All fluctuation theorems are statements that this Gibbs-weight is the source of thermodynamic irreversibility.

Structural Role

9.3 closes the loop opened by 9.1 and 9.2. It says: the Gibbs state (9.1), plugged into the TPM scheme (9.2), satisfies exact non-equilibrium identities. This has three consequences:

- The second law becomes a *theorem*, not a postulate, once microscopic reversibility and the Gibbs initial condition are granted.
- Free energy differences ΔF — usually equilibrium-only quantities — can be extracted from non-equilibrium measurements. This is the basis of experimental free-energy estimation (Liphardt et al. 2002 for classical; Batalhão et al. 2014 for quantum NMR).
- Landauer's principle (9.4) and engine bounds (9.5) are corollaries in special cases.

The quantum version is structurally identical to the classical one because the derivation only uses unitarity, TPM, and Gibbs weights — no commutativity assumption is needed.

Mathematical Development

Derivation of Jarzynski from TPM

Start from the definition:

$$\langle e^{-\beta W}\rangle_F = \sum_{n,m} p_n^0\, p_{m|n}\, e^{-\beta(E_m(\tau)-E_n(0))}.$$

Substitute $p_n^0 = e^{-\beta E_n(0)}/Z_0$:

$$= \frac{1}{Z_0}\sum_{n,m} e^{-\beta E_m(\tau)}\, p_{m|n}.$$

The quantum transition probabilities $p_{m|n} = |\langle m_\tau|U_\tau|n_0\rangle|^2$ satisfy $\sum_n p_{m|n} = 1$ (rows of a doubly stochastic matrix, by unitarity of U_τ). Therefore

$$\langle e^{-\beta W}\rangle_F = \frac{1}{Z_0}\sum_m e^{-\beta E_m(\tau)} = \frac{Z_\tau}{Z_0} = e^{-\beta\Delta F}. \qquad \blacksquare$$

The derivation uses three ingredients: (i) Gibbs initial condition, (ii) unitary forward evolution (so $p_{m|n}$ is doubly stochastic), (iii) TPM definition of W. No assumption about the protocol being slow, weak, or near-equilibrium enters.

Derivation of Crooks

Compute the joint forward probability:

$$P_F(n\to m) = \frac{e^{-\beta E_n(0)}}{Z_0}\,|\langle m_\tau|U_\tau|n_0\rangle|^2.$$

For the time-reversed protocol, start in $e^{-\beta H(\tau)}/Z_\tau$, apply $\tilde{U}_\tau = \Theta U_\tau^\dagger \Theta^{-1}$, and measure at the endpoints. Microreversibility of unitary dynamics gives

$$|\langle \tilde{n}_0|\tilde{U}_\tau|\tilde{m}_\tau\rangle|^2 = |\langle m_\tau|U_\tau|n_0\rangle|^2.$$

Hence

$$P_R(m \to n) = \frac{e^{-\beta E_m(\tau)}}{Z_\tau} \, |\langle m_\tau | U_\tau | n_0 \rangle|^2.$$

Divide the two:

$$\frac{P_F(n \to m)}{P_R(m \to n)} = \frac{Z_\tau}{Z_0} \, e^{-\beta(E_n(0) - E_m(\tau))} = e^{\beta(W - \Delta F)}.$$

The left-hand ratio, summed over pairs with fixed W, gives $P_F(W)/P_R(-W)$. ■

Second law via Jensen

By Jensen's inequality on the convex function e^{-x}:

$$e^{-\beta\langle W\rangle} \le \langle e^{-\beta W}\rangle = e^{-\beta \Delta F} \quad \Longrightarrow \quad \langle W\rangle \ge \Delta F.$$

Jarzynski contains the second law as a one-line corollary. Any protocol for which $\langle W\rangle = \Delta F$ exactly must have a delta-peaked $P(W)$, i.e. be reversible.

Fluctuation-dissipation at linear response

Expanding Crooks around equilibrium $W \approx \Delta F$, the fluctuation theorem reproduces the linear-response fluctuation-dissipation theorem: the dissipated work $W - \Delta F$ has Gaussian statistics with variance $2k_B T\langle W - \Delta F\rangle$. Quantum quenches that stay near-equilibrium obey this Gaussian limit; far-from-equilibrium quenches exhibit the full non-Gaussian Crooks identity.

Worked Example

Sudden quench of a qubit

Take the sudden-quench protocol from 9.2: $H(0) = \frac{\omega}{2} Z$, $H(\tau) = \frac{\omega}{2}(\cos\theta Z + \sin\theta X)$, $U_\tau = I$. The forward work distribution is (from 9.2):

$$P_F(W=0) = \cos^2(\theta/2), \quad P_F(W=+\omega) = p_0^0 \sin^2(\theta/2), \quad P_F(W=-\omega) = p_1^0 \sin^2(\theta/2).$$

Check Jarzynski. Compute $\langle e^{-\beta W}\rangle_F$:

$$\cos^2(\theta/2) + \frac{\sin^2(\theta/2)}{1+e^{-\beta\omega}}(e^{-\beta\omega} + e^{-\beta\omega} e^{\beta\omega}) = \cos^2(\theta/2) + \sin^2(\theta/2) \cdot \frac{1+e^{-\beta\omega}}{1+e^{-\beta\omega}} = 1.$$

Compute $e^{-\beta\Delta F} = Z_\tau / Z_0$. Both Hamiltonians have eigenvalues $\pm\omega/2$, so $Z_\tau = Z_0$ and $\Delta F = 0$. Consistent: $\langle e^{-\beta W}\rangle_F = 1 = e^{-0}$. ✓

Check Crooks. The reverse protocol starts in $e^{-\beta H(\tau)}/Z_\tau$ and quenches back to $H(0)$. Working it out, one finds

$$P_R(W=0) = \cos^2(\theta/2), \quad P_R(W=+\omega) = p_0^\tau \sin^2(\theta/2), \quad P_R(W=-\omega) = p_1^\tau \sin^2(\theta/2),$$

where $p_0^\tau = 1/(1+e^{-\beta\omega}) = p_0^0$ because the new Hamiltonian has the same spectrum. The Crooks ratio at $W = +\omega$:

$$\frac{P_F(+\omega)}{P_R(-\omega)} = \frac{p_0^0 \sin^2(\theta/2)}{p_1^0 \sin^2(\theta/2)} = \frac{p_0^0}{p_1^0} = e^{\beta\omega} = e^{\beta(W-\Delta F)}. \checkmark$$

Experimental verification. This exact protocol (or its analogue) has been run in the lab: Batalhão et al. (2014) verified quantum Jarzynski for a chloroform NMR system, and An et al. (2015) did so for a trapped ion. Measured $P(W)$ agreed with theory within error bars across the full distribution, not just the mean.

Cross-Links

- **9.1 Gibbs state** — the initial condition whose Boltzmann weights drive the whole derivation.
- **9.2 Quantum work and heat** — supplies the TPM definition used in both Jarzynski and Crooks.
- **9.4 Landauer's principle** — a corollary of the second law in the special case of information erasure.
- **9.5 Quantum heat engines** — fluctuation theorems bound the efficiency fluctuations of finite-time cycles.
- **5.1 Density matrices / time reversal** — microreversibility $|\langle m_\tau | U_\tau | n_0 \rangle|^2 = |\langle \tilde{n}_0 | \tilde{U}_\tau | \tilde{m}_\tau \rangle|^2$.
- **Statistics** — Crooks is a detailed-balance relation between forward and reverse path measures.

Structural Compression

Invariant: $\langle e^{-\beta W} \rangle_F = e^{-\beta \Delta F}$ for any protocol taking a quantum system out of thermal equilibrium.

Mechanism: Gibbs initial state $\times$ double stochasticity of unitary transition probabilities $\times$ TPM definition of work.

Cross-System Echo: Classical Jarzynski (1997) and Crooks (1999) — derived from Hamiltonian dynamics + Gibbs. Quantum version derived from unitary dynamics + Gibbs. Same structural skeleton; quantum version is actually *easier* to prove because unitarity gives double stochasticity for free.

Exercises

1. Fill in the three lines of the Jarzynski derivation, identifying exactly where unitarity of U_τ is used and what goes wrong if U_τ is replaced by a non-unital CPTP map.

2. Derive Crooks from the joint ratio $P_F(n \to m)/P_R(m \to n)$ and verify that integrating over W reproduces Jarzynski.

3. Show that for an *isothermal reversible* protocol (infinitely slow quasi-static), $P(W)$ is a delta function and $\langle W \rangle = \Delta F$ exactly.

4. For a sudden quench of a harmonic oscillator, $H(0) = \hbar\omega(a^\dagger a + 1/2) \to \hbar\omega'(a^\dagger a + 1/2)$, compute the full work distribution and verify Jarzynski using the partition functions $Z = (2\sinh(\beta\hbar\omega/2))^{-1}$.

5. Show that the Kullback–Leibler divergence $D(P_F \| \tilde{P}_R)$, where $\tilde{P}_R(W) = P_R(-W)e^{-\beta W}/\langle e^{-\beta W} \rangle$, equals the dissipated work $\beta(\langle W \rangle - \Delta F)$. Interpret structurally.

6. Suppose an experiment reports $\langle W \rangle - \Delta F = 0.5\, k_B T$ with Gaussian statistics. Use the linear-response limit of Crooks to predict the variance, and verify consistency with a direct calculation.

7. **Argue, structurally**, why fluctuation theorems take exactly the same algebraic form in classical and quantum thermodynamics, despite the fact that classical work is a path integral and quantum work is a pair of measurement outcomes. Identify what feature of the derivation is *not* used in either case.

Landauer's Principle

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.4

Landauer's principle is the structural bridge between information theory and thermodynamics: any logically irreversible operation — paradigmatically, erasing a bit — has a non-zero minimum thermodynamic cost, $k_B T \log 2$ of heat per bit dissipated into the environment. It is the reason Maxwell's demon does not violate the second law, and the reason classical computation has a fundamental energetic floor. In the quantum setting it follows directly from the strong subadditivity of von Neumann entropy (7.4) together with the Gibbs-state fixed-point property (9.1).

Formal Statement

Landauer's principle (Landauer 1961; quantum formulation: Reeb–Wolf 2014): Let a memory M (system of interest) be reset by a protocol that couples it to a thermal bath B at inverse temperature $\beta = 1/(k_B T)$. Let the joint evolution be unitary on MB, with B initially in the Gibbs state $\rho_B^{\text{th}} = e^{-\beta H_B}/Z_B$ and uncorrelated with M. Then the average heat dissipated into the bath,

$$\langle Q \rangle = \operatorname{tr}\big((\rho'_B - \rho_B^{\text{th}})H_B\big)\,,$$

is bounded below by the entropy decrease of the memory:

$$\boxed{\beta\langle Q\rangle \geq S(\rho_M) - S(\rho'_M)}$$

where ρ_M, ρ'_M are the initial and final reduced memory states and S is von Neumann entropy. In particular, if the protocol resets the memory from a maximally mixed 1-bit state $\rho_M = I/2$ to a pure state $\rho'_M = |0\rangle\langle 0|$,

$$\langle Q \rangle \geq k_B T \log 2.$$

The Reeb–Wolf refinement replaces $\geq$ with a finite-size equality plus non-negative correction terms from residual MB correlations and bath non-thermality after the protocol. Equality with the Landauer bound is reached only in the thermodynamic limit with quasi-static erasure.

Intuition

A "bit" is a memory whose state-space has two distinguishable configurations. Erasing it means *resetting* all configurations to a single one, say $|0\rangle$. This is an entropy-decreasing operation on the memory — from 1 bit of entropy down to 0. The second law says total entropy can only increase, so something else must absorb at least 1 bit. The only available sink is the thermal bath, and the thermodynamic "currency" for 1 bit of entropy at temperature T is $k_B T \log 2$ joules.

The *why* is structural:

1. The memory's entropy can only decrease if the bath's entropy increases.
2. The bath's entropy can only increase if it absorbs heat.
3. The minimum heat required for a given entropy increase is set by $dS = dQ/T$ at the bath temperature.

Combining these three gives Landauer. It is not a statement about any particular physical implementation — it is a statement about what any physically realizable erasure *must* do to the total entropy. Maxwell's demon fails because the demon's memory must eventually be erased, and that erasure cost is exactly what the demon "saved" earlier.

Structural Role

Landauer's principle anchors the identification of Shannon entropy (bits) with thermodynamic entropy (k_B per nat). Without it the two would be formally analogous but dimensionally disconnected. With it they become physically the same quantity, up to a unit conversion $k_B \log 2$ per bit.

Downstream:

- **9.5 Quantum heat engines** inherit Landauer as a lower bound on any cycle that resets a memory register.
- **9.6 Resource theories** of thermodynamics count "pure states" as thermodynamic resources whose exchange rate with heat is $k_B T \log 2$ per bit.
- **Reversible computation** (Bennett 1973) is the logical converse: a computation without bit erasure has no Landauer cost, only finite-time dissipation.
- **Black-hole thermodynamics** uses a Landauer-like identification when assigning entropy $k_B \log 2$ per bit to horizon area.

Mathematical Development

Proof via second law on memory + bath

Let the joint state evolve unitarily: $\rho'_{MB} = U\rho_M \otimes \rho_B^{\text{th}} U^\dagger$. Unitary evolution preserves total entropy:

$$S(\rho'_{MB}) = S(\rho_{MB}) = S(\rho_M) + S(\rho_B^{\text{th}}).$$

By subadditivity of von Neumann entropy (7.4),

$$S(\rho'_{MB}) \leq S(\rho'_M) + S(\rho'_B).$$

Combine:

$$S(\rho_M) + S(\rho_B^{\text{th}}) \leq S(\rho'_M) + S(\rho'_B).$$

Rearrange:

$$S(\rho_M) - S(\rho'_M) \leq S(\rho'_B) - S(\rho_B^{\text{th}}).$$

Now use the fact that the Gibbs state maximizes entropy at fixed expected energy. Concretely, Klein's inequality applied to ρ'_B and ρ_B^{th} gives

$$S(\rho'_B) - S(\rho_B^{\text{th}}) \leq \beta \operatorname{tr}\big((\rho'_B - \rho_B^{\text{th}})H_B\big) = \beta\langle Q\rangle.$$

Chaining the two inequalities:

$$S(\rho_M) - S(\rho'_M) \leq \beta\langle Q\rangle. \qquad \blacksquare$$

Where each step is used

- **Unitarity of joint evolution**: ensures total entropy is conserved, so memory-entropy decrease = bath-entropy increase.
- **Subadditivity**: accounts for any residual correlations between M and B after the protocol. Equality requires the post-protocol state to be product (uncorrelated memory + bath).
- **Gibbs being the max-entropy state**: ensures heat-to-entropy conversion at the bath achieves at best the rate $1/k_BT$. Equality requires the bath to remain thermal after erasure.

Both equality conditions can only be met in the idealized quasi-static, thermodynamic-limit case. Real erasures have a finite-size Reeb–Wolf correction.

Reeb–Wolf equality

Reeb and Wolf (2014) proved the exact version:

$$\beta\langle Q\rangle = S(\rho_M) - S(\rho'_M) + I(M' : B') + D(\rho'_B \| \rho_B^{\text{th}}),$$

where $I(M' : B')$ is the mutual information between memory and bath after the protocol and D is the relative entropy. Both correction terms are non-negative, so Landauer's bound is recovered, and it shows explicitly what needs to vanish for equality.

Worked Example

Erasing a single qubit memory

Let M be a single qubit with trivial internal Hamiltonian (degenerate), initially in $\rho_M = I/2$ (maximally mixed, 1 bit of entropy). The bath is a large thermal reservoir at temperature T. The erasure protocol:

1. **Couple** M to an auxiliary qubit with Hamiltonian $H_M = \epsilon|1\rangle\langle 1|$, adiabatically raised from $\epsilon = 0$ to $\epsilon \gg k_BT$. Under slow coupling + bath contact, M equilibrates at each ϵ to $e^{-\beta H_M}/Z$.
2. At $\epsilon \gg k_BT$, the Gibbs state is $\approx |0\rangle\langle 0|$ — the memory is erased.
3. **Decouple** by lowering $\epsilon \to 0$ again. If done slowly, the state remains $|0\rangle\langle 0|$.

Work done on M during the ramp-up (reversible, integrated over ϵ):

$$W_{\text{ramp}} = \int_0^\infty \partial_\epsilon F(\epsilon)\, d\epsilon = \int_0^\infty \frac{\epsilon\, e^{-\beta\epsilon}}{1 + e^{-\beta\epsilon}}\, d\epsilon \cdot \beta\text{-factor},$$

which evaluates to exactly $k_BT \log 2$ in the quasi-static limit.

Heat dissipated to the bath: $\langle Q\rangle = k_BT \log 2$ (by the first law, since $\Delta U_M = 0$ for the cycle $0 \to \epsilon \to 0$).

Numerical value at $T = 300$ K:

$$k_BT \log 2 \approx (1.38 \times 10^{-23}\,\text{J/K})(300\,\text{K})(0.693) \approx 2.87 \times 10^{-21}\,\text{J} \approx 17.9\,\text{meV}.$$

Modern CMOS gates dissipate $\sim 10^{-18}$ J per bit operation, still three orders of magnitude above Landauer's bound. Recent experiments (Bérut et al. 2012 classical; Yan et al. 2018 quantum Josephson) have directly measured heat flow during bit erasure and verified the bound to within the Reeb–Wolf correction terms.

Cross-Links

- **7.4 Von Neumann entropy** — supplies subadditivity $S(\rho_{MB}) \leq S(\rho_M) + S(\rho_B)$ and Klein's inequality.
- **9.1 Gibbs state** — the max-entropy characterization gives the $\leq \beta\langle Q\rangle$ step.
- **9.2 Work and heat** — identifies what $\langle Q\rangle$ *is* in the TPM scheme.
- **9.5 Quantum heat engines** — any cycle that includes a reset step pays a Landauer cost.
- **Information theory** — identifies Shannon bit with physical bit; Maxwell's demon resolution (Bennett 1982).
- **Reversible computation** — Bennett/Fredkin/Toffoli gates avoid erasure and thus Landauer cost.

Structural Compression

Invariant: Erasing ΔS nats of memory entropy requires dissipating at least $\Delta S/\beta$ joules of heat into the bath.

Mechanism: Subadditivity of joint entropy + Klein's inequality against the Gibbs state.

Cross-System Echo: Same form in classical Landauer (derived from Shannon entropy) and quantum Landauer (derived from von Neumann entropy). The structural identity $S_{\text{info}} \leftrightarrow S_{\text{thermo}}$ is the content of the theorem; everything else is bookkeeping.

Exercises

1. Prove that if the joint post-protocol state ρ'_{MB} is uncorrelated ($I(M' : B') = 0$) and the bath remains thermal ($D(\rho'_B \| \rho_B^{\text{th}}) = 0$), then Landauer's inequality is saturated. Interpret each condition physically.

2. Compute the Landauer cost of erasing an n-bit memory that starts in a state of Shannon entropy H bits and ends in the all-zero state. Show the answer is $n \log 2 - H$ in units of $k_B T$, *not* $n \log 2$.

3. Derive Klein's inequality $S(\rho) - S(\sigma) \leq \operatorname{tr}(\rho - \sigma) \log \sigma^{-1}$ from the non-negativity of quantum relative entropy, and use it to justify the $\beta\langle Q\rangle$ step in the proof.

4. A Maxwell demon observes a single-molecule gas and records which half of the container the molecule is in, then uses this to extract $k_B T \log 2$ of work. Show that the demon's memory must hold 1 bit, and that erasing this memory at the end of the cycle costs exactly $k_B T \log 2$, restoring the second law.

5. A quantum erasure protocol takes τ seconds and dissipates $k_B T \log 2 + \Delta$ per bit, with $\Delta > 0$ due to finite-time effects. Show that, to leading order in $1/\tau$, $\Delta \propto 1/\tau$ and identify the prefactor in terms of the relaxation time of the bath.

6. For a memory with non-degenerate Hamiltonian H_M, show that the Landauer bound generalizes to $\beta\langle Q\rangle \geq S(\rho_M) - S(\rho'_M) - \beta\Delta U_M$, where $\Delta U_M = \operatorname{tr}(\rho'_M H_M) - \operatorname{tr}(\rho_M H_M)$. Interpret the new term as work done on the memory.

7. **Argue, structurally**, why Landauer's principle is the *unique* way to reconcile reversible microdynamics (unitary quantum mechanics) with irreversible logical operations (bit erasure) while preserving the second law. What would break if the coefficient $k_B T \log 2$ were replaced by anything smaller?

Quantum Heat Engines

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.5

A quantum heat engine extracts work from temperature differences using a quantum system as its working substance. The cycle structure is the same as a classical engine — isothermal contact with a hot bath, expansion, isothermal contact with a cold bath, compression — but the expansions and compressions are protocol-driven changes of a *quantum* Hamiltonian and the equilibrations use a *Lindblad* thermalization from 8.2. The question of node 9.5 is: do quantum effects change the efficiency? The structural answer is **no at the asymptotic Carnot limit, yes at finite-time power and in specific non-equilibrium regimes**.

Formal Definition

A **quantum heat engine** is a cyclic protocol on a quantum system S with a parameter-dependent Hamiltonian $H(\lambda)$, alternately coupled to a hot bath at temperature T_h and a cold bath at temperature T_c. Over one cycle, the state returns to itself: $\rho_{\text{end}} = \rho_{\text{start}}$. By the first law applied to a cycle,

$$\oint d\langle E\rangle = 0 \quad \Longrightarrow \quad W_{\text{net}} + Q_h + Q_c = 0,$$

where $Q_h > 0$ is heat absorbed from the hot bath, $Q_c < 0$ is heat dumped to the cold bath, and $W_{\text{net}} < 0$ is work extracted (with the sign convention that work done *on* the system is positive).

Efficiency is defined as

$$\eta = \frac{|W_{\text{net}}|}{Q_h} = 1 - \frac{|Q_c|}{Q_h}.$$

Carnot bound: For any cyclic engine operating between $T_h > T_c$,

$$\eta \leq \eta_C = 1 - \frac{T_c}{T_h}.$$

Equality holds only for reversible (quasi-static) cycles. This bound follows from the Clausius form of the second law $\oint dQ/T \leq 0$, which survives unchanged in the quantum setting because it follows from the Gibbs fixed-point property of Lindblad thermalization (9.1 + 8.2).

Intuition

Think of a working substance with a tunable energy gap ω. When you couple it to the hot bath, it fills up with thermal populations characteristic of T_h, ω_h. When you decouple it and compress (raise ω), you do work on it without changing populations — its internal energy goes up. When you couple to the cold bath, it sheds heat until equilibrated at T_c, ω_h. Finally, decouple and expand (lower ω) to close the cycle, extracting work.

The net effect over a full cycle: heat flows from hot to cold, and a fraction $\eta \leq \eta_C$ of the hot-side heat comes out as work. The "quantumness" enters in:

- The working substance is quantized — it's a qubit or oscillator, not a classical gas.
- Between strokes, the system can be in superpositions or have coherences.
- Coherences can be exploited (squeezed-bath engines, coherence-enhanced engines) but only within the Carnot constraint imposed by the second law.

There is no such thing as a quantum engine that beats Carnot. If a paper claims otherwise, it is tacitly redefining one of T_h, T_c, W, or Q.

Structural Role

9.5 is the operational arena for testing quantum thermodynamic concepts in the lab. Every theoretical tool from this module is used:

- The **Gibbs state** (9.1) sets the equilibrium endpoints.
- The **TPM work distribution** (9.2) gives statistics per stroke.
- **Fluctuation theorems** (9.3) constrain finite-time dissipation.
- **Landauer's principle** (9.4) applies if the cycle involves bit reset.
- **Lindblad thermalization** (8.2) models the bath contact strokes.

Quantum heat engines have been built and run: from trapped-ion heat engines (Roßnagel et al. 2016) to NMR engines (Peterson et al. 2019) to superconducting-circuit heat engines. Efficiencies up to ~30% of Carnot have been measured; the deviation is entirely due to finite-time dissipation, not any violation of the Carnot bound.

Mathematical Development

Quantum Otto cycle on a qubit

The Otto cycle has four strokes. Let the qubit have $H(\lambda) = \frac{\lambda}{2} Z$, with λ tunable between ω_c and ω_h, and $\omega_h > \omega_c$.

Stroke 1: Isochoric hot ($\lambda = \omega_h$, coupled to T_h). Start in ρ_A, end in $\rho_B = e^{-\beta_h \omega_h Z/2}/Z_{h,B}$. Heat absorbed:

$$Q_1 = \mathrm{tr}(\rho_B H_h) - \mathrm{tr}(\rho_A H_h).$$

Stroke 2: Adiabatic compression/expansion ($\lambda : \omega_h \to \omega_c$, decoupled from bath). The state follows U_{ad}. For the qubit, adiabaticity means populations in $H(\lambda)$ eigenstates are preserved. No heat, only work:

$$W_2 = \mathrm{tr}(\rho_B H_c) - \mathrm{tr}(\rho_B H_h).$$

Since populations don't change and the energy gap shrinks ($\omega_h \to \omega_c$), $W_2 < 0$ (work out).

Stroke 3: Isochoric cold ($\lambda = \omega_c$, coupled to T_c). End in $\rho_D = e^{-\beta_c \omega_c Z/2}/Z_{c,D}$. Heat released to cold bath:

$$Q_3 = \mathrm{tr}(\rho_D H_c) - \mathrm{tr}(\rho_B H_c) < 0.$$

Stroke 4: Adiabatic compression ($\lambda : \omega_c \to \omega_h$).

$$W_4 = \mathrm{tr}(\rho_D H_h) - \mathrm{tr}(\rho_D H_c) > 0.$$

By cycle closure, $\rho_D \to \rho_A$ under the compression, so the equilibration at T_h restores ρ_B with the same populations as ρ_A carried over from ρ_D plus bath relaxation.

Otto efficiency

The quantum Otto efficiency for a qubit between (T_h, ω_h) and (T_c, ω_c) is

$$\eta_{\mathrm{Otto}} = 1 - \frac{\omega_c}{\omega_h}.$$

Compare with Carnot $\eta_C = 1 - T_c/T_h$. The Otto cycle reaches Carnot only when

$$\frac{\omega_c}{\omega_h} = \frac{T_c}{T_h},$$

i.e. when the compression ratio matches the temperature ratio. Otherwise $\eta_{\text{Otto}} < \eta_C$, and the engine is strictly sub-Carnot even though its specific efficiency formula looks "simpler."

At the Carnot limit the cycle reduces to zero net power (infinite cycle time). The curve-Novikov-Chambadal bound at **maximum power** is $\eta_{\text{MP}} = 1 - \sqrt{T_c/T_h}$, and this is the more relevant figure of merit for real devices.

Why the Carnot bound survives quantization

The Carnot bound comes from applying the Clausius inequality $\oint dQ/T \leq 0$ to the cycle. In the quantum case this becomes

$$\sum_{\text{baths}} \int \frac{d\langle Q_{\text{bath}}\rangle}{T_{\text{bath}}} \leq 0,$$

which follows from strong subadditivity of von Neumann entropy + the Gibbs state being the max-entropy state. There is no step in the derivation that uses commutativity, classicality, or absence of superposition. Hence the bound survives.

Worked Example

Trapped-ion Otto engine (Roßnagel et al. 2016 style)

A single $^{40}\text{Ca}^+$ ion in a tapered Paul trap acts as the working substance. The "gap" is the spacing of vibrational levels, tuned by the trap frequency $\nu(t)$ between $\nu_h = 1$ MHz and $\nu_c = 200$ kHz. Hot bath: electric-field noise injected at $T_h \sim 10^4$ K effective temperature. Cold bath: spontaneous emission at $T_c \sim 10^2$ K.

Temperature ratio: $T_c/T_h = 0.01$, so $\eta_C \approx 0.99$.

Frequency ratio: $\nu_c/\nu_h = 0.2$, so $\eta_{\text{Otto}} = 0.8$.

Measured efficiency: ~0.28, limited by finite-time dissipation, phonon leakage, and imperfect thermalization. The measurement verified the structural prediction: the cycle performs as a heat engine, efficiency obeys the Otto formula as an upper bound, and Carnot is never approached because the cycle is fast relative to the relaxation time.

Interpretation in the TPM framework: work per stroke is a random variable; only the average obeys the clean formulas above. Individual runs exhibit the full Crooks distribution.

Zero-power Carnot limit

Take $\omega_h \to \omega_c$ at fixed T_h, T_c. The gap difference vanishes, so the adiabatic strokes do zero work per cycle. But the isothermal strokes also transfer zero net heat because the populations at T_h, ω and T_c, ω differ only by a factor $\sim (\beta_h - \beta_c)\omega$. The ratio $\eta \to 1 - T_c/T_h$ is recovered by quasi-static integration. This illustrates the general rule: high efficiency $\leftrightarrow$ low power.

Cross-Links

- **9.1 Gibbs state** — thermalization endpoints of each isothermal stroke.
- **9.2 TPM work** — stochastic work per stroke.
- **9.3 Fluctuation theorems** — constrain dissipation in finite-time cycles.
- **9.4 Landauer** — applies if the engine includes any memory reset.
- **8.2 Lindblad thermalization** — models the bath-contact strokes.
- **5.1 Density matrices** — the engine state is generally mixed during the cycle.
- **Classical thermodynamics** — Otto, Carnot, Stirling cycles all have direct quantum analogues.

Structural Compression

Invariant: $\eta \le \eta_C = 1 - T_c/T_h$ for any cyclic engine between two thermal baths, quantum or classical.

Mechanism: Clausius inequality applied via strong subadditivity + Gibbs max-entropy.

Cross-System Echo: Classical Carnot (1824). Quantum Otto, quantum Carnot, quantum Stirling — all obey the same bound. Quantum coherence can shift *power*, not *efficiency limit*.

Exercises

1. Derive the Otto efficiency $\eta_{\text{Otto}} = 1 - \omega_c/\omega_h$ for a qubit working substance, starting from the per-stroke heat calculations.
2. Show that a quantum Carnot cycle (two isothermals + two adiabatics, all quasi-static) reaches η_C in the quasi-static limit, and that the net power vanishes as τ^{-1}.
3. For a quantum harmonic-oscillator Otto cycle with $H(\lambda) = \hbar\lambda(a^\dagger a + 1/2)$, compute Q_h, Q_c, W_{net} in terms of $\lambda_h, \lambda_c, T_h, T_c$, and verify $\eta = 1 - \lambda_c/\lambda_h$.
4. Show that a coherence-enhanced engine (starting the stroke in a superposition of energy eigenstates) has the same average work output as a diagonal one *to first order*, but non-zero work *variance*. Relate this to the linear-response limit of Crooks.
5. Compute the Curzon–Ahlborn efficiency at maximum power $\eta_{\text{MP}} = 1 - \sqrt{T_c/T_h}$ for a quantum Otto engine in the high-temperature limit, and show that it is always less than η_C but can exceed $\eta_C/2$.
6. A squeezed-bath engine uses a bath whose state is $e^{-\beta H_{\text{bath}}}/Z$ squeezed by a Bogoliubov transformation. Show that it can *appear* to exceed Carnot if one misidentifies the effective temperature, and correct the analysis by computing the actual entropy production.
7. **Argue, structurally**, why the Carnot bound is the same in classical and quantum thermodynamics. Identify the step in the derivation where classical commutativity *could* in principle matter, and explain why it does not.

Resource Theories

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Quantum Thermodynamics **Node:** 9.6

A **resource theory** is a framework in which certain states and certain operations are declared "free," and everything else is quantified as a costly resource relative to that choice. Quantum thermodynamics becomes a resource theory when the free states are Gibbs states and the free operations are *thermal operations*. The resulting framework generalizes the second law from a single inequality into an infinite family of monotones — one for each Rényi divergence — and gives an exact, non-asymptotic characterization of which non-equilibrium states can be converted into which. This is the modern foundation of quantum thermodynamics (Brandão, Horodecki, Oppenheim, Renes, Spekkens, and others, mid-2010s).

Formal Definition

A **resource theory** is specified by three ingredients:

1. **Free states** $\mathcal{F} \subset \mathcal{S}(\mathcal{H})$: the states you are allowed to have at no cost.
2. **Free operations** $\mathcal{O}$: the CPTP maps you are allowed to apply at no cost. $\mathcal{O}$ must map $\mathcal{F} \to \mathcal{F}$.
3. **Resource monotones** $M : \mathcal{S}(\mathcal{H}) \to \mathbb{R}_{\geq 0}$: functions that are non-increasing under free operations. $M(\rho) \leq M(\Lambda(\rho))\dots$ sorry, $M(\Lambda(\rho)) \leq M(\rho)$ for all $\Lambda \in \mathcal{O}$.

Thermal resource theory (with respect to Hamiltonian H, temperature T):

- **Free states:** The Gibbs state $\rho_\beta = e^{-\beta H}/Z$ is the *only* free state for each subsystem. (Any other state is a resource called "athermality.")
- **Free operations (thermal operations):** CPTP maps Λ on the system S that admit a Stinespring dilation of the form
$$\Lambda(\rho_S) = \mathrm{tr}_B\left[U(\rho_S \otimes \rho_B^{\mathrm{th}})U^\dagger\right],$$
where $\rho_B^{\mathrm{th}} = e^{-\beta H_B}/Z_B$ is the bath's own Gibbs state and U commutes with the total Hamiltonian $H_S + H_B$ (strict energy conservation).
- **Monotones:** The generalized free energies $F_\alpha(\rho) = \beta^{-1}[S_\alpha(\rho\|\rho_\beta) - \log Z]$ for each Rényi order $\alpha \geq 0$, where $S_\alpha(\rho\|\sigma) = (\alpha-1)^{-1}\log\mathrm{tr}(\rho^\alpha\sigma^{1-\alpha})$. The standard free energy is the $\alpha = 1$ case.

The main theorem: a state ρ can be converted to σ by thermal operations if and only if $F_\alpha(\rho) \geq F_\alpha(\sigma)$ for **all** $\alpha \geq 0$ (quasi-classical case), or equivalently ρ *thermomajorizes* σ.

Intuition

The ordinary second law says "the free energy $F = U - TS$ can only decrease under thermal operations." That gives you *one* inequality. For asymptotically many copies (thermodynamic limit) it is tight: if $F(\rho) > F(\sigma)$, you can convert many copies of ρ into many copies of σ. But for *single shot* or *finite-copy* transformations, one inequality is too weak. You need an entire family, one per Rényi order, each giving a different constraint on the same transformation.

The reason is structural: the Gibbs state is a complicated object with many "moments" of resourcefulness, and the single-copy free energy only captures the first moment. The higher Rényi divergences capture higher moments (how resource-rich is the state in its tail?), and each generates its own second law. In the asymptotic limit all Rényi divergences converge to the standard relative entropy, and you recover the usual one-line second law.

This is the central structural shift of 21st-century quantum thermodynamics: the second law is not one inequality but a *family*.

Structural Role

9.6 recasts the entire module in a clean, high-level language:

- **9.1 Gibbs state** = the unique free state.
- **9.2 Work** = the athermality of a pure energy-eigenstate resource (the "work bit" $|w\rangle$).
- **9.3 Fluctuation theorems** = consequences of contractivity of relative entropy under thermal operations.
- **9.4 Landauer** = a monotone statement about the $\alpha = 1$ free energy for the erasure transformation.
- **9.5 Heat engines** = cycles in which work is input and output, and the Carnot bound is one inequality from the full family.

It also connects quantum thermodynamics to *other* resource theories:

- **Entanglement** (free = separable states, free = LOCC; monotones = entanglement measures).
- **Coherence** (free = incoherent states, free = incoherent operations; monotones = coherence measures).
- **Magic** (free = stabilizer states, free = Clifford gates; monotones = magic measures).
- **Asymmetry** (free = symmetric states, free = covariant operations).

The resource-theoretic template is the same in all of them. Quantum thermodynamics is one instance among several.

Mathematical Development

Thermomajorization

For a classical (diagonal) state with eigenvalues $p_1, ..., p_d$ at energies $E_1, ..., E_d$, the **thermomajorization curve** is the piecewise-linear function obtained by:

1. Sort the levels by the β-exponential weight $e^{\beta E_i} p_i$ in decreasing order.
2. Plot the cumulative sum of p_i against the cumulative sum of $e^{-\beta E_i}$ on the x-axis.

A state p **thermomajorizes** q, written $p \succ_\beta q$, if its thermomajorization curve lies everywhere on or above q's.

Horodecki–Oppenheim theorem (2013): For quasi-classical states, ρ can be converted to σ by thermal operations if and only if $\rho \succ_\beta \sigma$.

This is the quantum generalization of Hardy–Littlewood–Polya's classical majorization theorem (which recovers the case $\beta = 0$, i.e. $H = 0$).

Generalized free energies

For each Rényi order $\alpha \geq 0$ define

$$F_\alpha(\rho) = \frac{1}{\beta}\left[S_\alpha(\rho \| \rho_\beta) - \log Z\right], \quad S_\alpha(\rho\|\sigma) = \frac{1}{\alpha - 1} \log \operatorname{tr}(\rho^\alpha \sigma^{1-\alpha}).$$

Special cases:

- $\alpha = 0$: F_0 — counts the rank of ρ.
- $\alpha = 1$: $F_1 = \operatorname{tr}(\rho H) - TS(\rho)$ — the standard Helmholtz free energy.
- $\alpha = 2$: relates to the second moment / purity.
- $\alpha = \infty$: F_∞ — single-shot worst-case, associated with min-entropy.

Each F_α is monotonically non-increasing under thermal operations. For quasi-classical states, the conjunction of all these inequalities is exactly equivalent to thermomajorization.

Brandão et al. second laws (2015)

The full set of second laws for single-shot thermodynamics:

$$F_\alpha(\rho) \geq F_\alpha(\sigma) \quad \text{for all } \alpha \geq 0.$$

In the asymptotic i.i.d. limit (many copies, averaging), all F_α converge to F_1, and a single second law $F_1(\rho) \geq F_1(\sigma)$ suffices. Finite-copy thermodynamics is genuinely richer than the asymptotic second law.

Catalytic thermal operations

Allowing a "catalyst" — an auxiliary system that is returned unchanged — expands the set of reachable states. The classical result says catalysts do not give new constraints beyond thermomajorization, but they can smooth out the transformation. This is analogous to catalysts in classical chemistry.

Worked Example

Single-qubit thermomajorization

Take a qubit with $H = \frac{\omega}{2}(I - Z)$, so energies are 0 and ω. Gibbs populations at inverse temperature β:

$$\pi = (\pi_0, \pi_1) = \left(\frac{1}{1 + e^{-\beta\omega}}, \ \frac{e^{-\beta\omega}}{1 + e^{-\beta\omega}} \right).$$

Consider two diagonal states: $\rho = (0.8, 0.2)$ and $\sigma = (0.7, 0.3)$.

Compute weights: $e^{\beta E_i} p_i$ for ρ: $(0.8, 0.2e^{\beta\omega})$. If $e^{\beta\omega} > 4$, then $0.2e^{\beta\omega} > 0.8$ and the ordering is (excited, ground). Otherwise (ground, excited).

Case $\beta\omega = 0$ (infinite temperature, Gibbs is maximally mixed):

Thermomajorization reduces to standard majorization. The curve for ρ plots $(0,0) \to (0.5, 0.8) \to (1,1)$; for σ, $(0,0) \to (0.5, 0.7) \to (1,1)$. ρ lies above σ, so $\rho \succ_\beta \sigma$ and $\rho \to \sigma$ is allowed.

Case $\beta\omega = 2$: Need to reorder by $e^{\beta E_i} p_i$. $e^{\beta\omega} p_1 = e^2 \cdot 0.2 \approx 1.48 > 0.8$, so we order (level 1, level 0). Curve for ρ uses x-coordinates $(0, e^{-\beta\omega}, 1) \approx (0, 0.135, 1)$. Compute and compare. At the intermediate point, ρ still lies above σ, so the transformation is allowed.

Case ρ below Gibbs, σ above Gibbs: Transformation is forbidden because thermal operations cannot increase the free energy.

Case both are Gibbs: Thermomajorization is trivially saturated in both directions; identity is the only transformation.

Work extraction as a resource process

A pure excited state $|1\rangle\langle 1|$ is an athermal resource. Extracting work from it means converting $|1\rangle$ to $|0\rangle$ while raising a work-storage system from $|w_0\rangle$ to $|w_1\rangle$, where $H_W|w_i\rangle = w_i|w_i\rangle$. The maximum extractable work is

$$W_{\max} = k_B T \log\left(\frac{1}{\pi_1}\right) = k_B T \log(1 + e^{\beta\omega}),$$

which for $\beta\omega \gg 1$ gives $W_{\max} \approx \omega$, and for $\beta\omega \ll 1$ gives $W_{\max} \approx k_B T \log 2$ (Landauer-like).

Cross-Links

- **9.1 Gibbs state** — the unique free state of thermal resource theory.
- **9.2 Work** — formalized as the athermality of a pure energy-eigenstate resource.
- **9.3 Fluctuation theorems** — consequences of monotone contractivity.
- **9.4 Landauer** — the $\alpha = 1$ monotone for the erasure transformation.
- **4.x Entanglement** — parallel resource theory; free = separable states, free = LOCC.
- **Coherence, magic, asymmetry** — further examples of the same structural template.
- **Information theory** — Rényi divergences are the core object; monotonicity under CPTP maps is the "data processing inequality."

Structural Compression

Invariant: Thermal operations define a partial order on states via thermomajorization, equivalently by preservation of all generalized free energies F_α.

Mechanism: Energy-conserving unitaries on system + Gibbs bath + partial trace = thermal operations. Monotones are Rényi divergences against Gibbs.

Cross-System Echo: Same template for entanglement (LOCC), coherence (incoherent ops), magic (Clifford), asymmetry (covariant ops). Quantum thermodynamics is one resource theory among several.

Exercises

1. Show that thermal operations map the Gibbs state to itself, and that this is the defining property of a "free state" in the resource-theory sense.

2. Derive the standard free energy $F_1(\rho) = \mathrm{tr}(\rho H) - TS(\rho)$ as the $\alpha \to 1$ limit of $F_\alpha(\rho) = \beta^{-1}[S_\alpha(\rho\|\rho_\beta) - \log Z]$ using L'Hôpital on the Rényi divergence.

3. For a qubit with energies $0, \omega$ at inverse temperature β, compute F_0, F_1, F_∞ for the pure excited state $|1\rangle\langle 1|$ and verify that all three are positive.

4. Show that two states with the same ordered eigenvalue list but different energy assignments have different thermomajorization curves, so the partial order really does depend on the Hamiltonian.

5. A thermal operation on a qubit cannot create coherence between energy eigenstates. Prove this using strict energy conservation of the Stinespring dilation.

6. Show that the Horodecki–Oppenheim thermomajorization criterion reduces to standard majorization (Hardy–Littlewood–Polya) in the infinite-temperature limit $\beta \to 0$. Interpret the result.

7. **Argue, structurally**, why a single second law is insufficient for single-shot quantum thermodynamics, and why exactly the Rényi family is the right generalization (rather than, say, arbitrary other convex functionals).