

# Thermodynamics — Narrative

Date: September 27, 2025 • Version: Draft (DO NOT OVERWRITE prior versions)

Constants: CODATA-2022 (NIST SP 961, May 2024).  $k_B$  is an acceptance lock.

## Locks & Conventions

### No imports (what we don't do):

We don't assume ideal-gas, Navier–Stokes, Newton/Fourier/Fick, or mean-field templates. Those forms (when they appear) are dictionaries of our route-counting and free-energy derivatives—not axioms we start from.

- Single dimensional scale:  $S_0 = \hbar$  (SI-locked). No other tunable scales.
- Route/closure language: same as the other pillars (display-area  $A_d$ ; route action  $S[\text{path}] \equiv \int A_d ds$  when we need phase).
- Data anchors: use canonical EoS/transport datasets; prefer ratios when tighter than absolutes.
- Notation sanity: new symbols are dimensionless unless explicitly units-bearing; absolutes are anchored in Calibration.

### Scope (what we do and what we don't)

We turn the same closure/torsion/shear budgets into: (i) equation-of-state (EoS) relations, (ii) transport coefficients via linear response (Onsager/FDT), and (iii) phase behavior and stability. We introduce no new dimensional scales beyond  $S_0 = \hbar$ ; we import  $k_B$  as an acceptance lock. Equilibrium and near-equilibrium are covered. Far-from-equilibrium and device-level microphysics lives in branches.

When it turns negative, the cheapest arrangement is no longer well-mixed ( $m=0$ ); the system saves score by picking  $|m|>0$ . That's the new phase—no extra dial needed.

Plain English: we count how many route-arrangements fit the macroscopic totals, pick the cleanest arrangement (a minimum of the right free-energy), and read off the pressures, chemical potentials, heat capacities, and transport numbers as derivatives. No black-box models, no hand-fit curves.

### Coverage map

- Foundation —  $S_0 = \hbar$ ; route counting; free-energy levers ( $F, G$ ) and constraints.
- What are  $T$  and  $S$ ? — temperature as trade-strength between routes; entropy as the logarithm of compatible arrangements.
- Free-energy levers — minima decide phases; derivatives give EoS and response.
- From budgets to EoS — how closure/torsion/shear enter the free-energy density without new scales.

- Fluctuations & FDT — why noise equals response; susceptibilities and Einstein relations.
- Transport — linear response  $J = L \cdot X$ ; Onsager reciprocity; diffusion/viscosity/thermal conductivity scalings.
- Phase behavior — order parameters and onset; Landau-style expansions and stability.
- What flips at  $T_c$ : the sign of the quadratic piece  $a(T)$ .
- Bridges — to EM (media response) and Mechanics/GR (work/clock links).
- Calibration & falsifiers — ratio-first targets; explicit pass/fail gates.
- Reader map — where the math and data live.

### What are Temperature and Entropy here? (same ontology, plain talk)

Temperature is “how hard the system pushes to trade energy between compatible routes.” Entropy is the log of how many ways you can arrange those routes and still hit the macroscopic totals. Nothing mystical and nothing added—just our route picture with a counting layer you can differentiate when you need levers.

$$S = k_B \cdot \ln \Omega \quad (\Omega = \text{number of compatible micro-routes})$$

$$\beta = 1 / (k_B T) \quad (\text{multipliers weight routes by energy})$$

If you clamp the totals (energy, volume, particle count), the maximum-entropy arrangement fixes the weights on each route. That’s why the standard ensembles appear without being assumed—we simply read them off of the counting. A good gut-check is the ice-water beaker: at the melting point you have many ways to shuffle energy without moving the totals much, so entropy is big and trade is easy—i.e., temperature is set by how easily that trade happens.

### Free-energy levers (why minima decide phases)

Free energies are scorecards made to match the constraints you are holding fixed; their minima pick the state. Once you’ve picked the right scorecard, the useful numbers are just slopes and curvatures—pressures, chemical potentials, capacities—no extra dials appear mid-derivation.

$$F(T, V, N) = E - T S$$

$$G(T, p, N) = H - T S \equiv F + pV$$

$$P = - (\partial F / \partial V)_{\{T, N\}} \quad \mu = (\partial F / \partial N)_{\{T, V\}}$$

$$\kappa_T = - (1/V) (\partial V / \partial P)_{\{T\}} \quad C_V = (\partial E / \partial T)_{\{V\}} \quad C_P = (\partial H / \partial T)_{\{p\}}$$

Read these in plain language: pressure is how hard the score goes up if you try to squeeze the system; chemical potential is the ‘price’ of adding one more particle; heat capacities tell you how much energy it takes to move the temperature by one degree under the constraint you chose.

## From VMS budgets to EoS (how geometry enters numbers)

What the math is saying: nudge  $T$  or  $p \rightarrow \mathcal{T}, \mathcal{B}, \mathcal{C}$  shift a little  $\rightarrow$  the slopes of  $F = \int f dV$  turn those geometric nudges into  $P, \kappa_T, C_V, \dots$ . One clean example lands in the Appendix.

What the indices are:  $\mathcal{T}$  (torsion),  $\mathcal{B}$  (shear), and  $\mathcal{C}$  (closure) summarize micro-geometry the routes live in; the  $a_\bullet$  are dimensionless weights fixed in this pillar.

We don't guess an EoS. We write the free-energy density from the same geometric budgets we've used all along, with dimensionless coefficients named in this pillar, and differentiate. Calibration sets material constants; ratios between nearby states usually cancel those and give tighter checks.

$$f = f_0 + a_T \cdot \mathcal{T} + a_S \cdot \mathcal{B} + a_C \cdot \mathcal{C} + \dots \quad (\text{dimensionless } a_\bullet; \text{ no new scales})$$

Plain English:  $\mathcal{T}, \mathcal{B}$ , and  $\mathcal{C}$  are the torsion, shear, and closure indices that the routes live in. Changing state (say, by nudging  $T$  or  $p$ ) changes those indices a little; the derivatives of  $F = \int f dV$  translate that geometry change into pressures and moduli.

## Fluctuations and FDT (why noise equals response)

Every measurable 'wiggle' around equilibrium comes from the same counting. The variance of a quantity  $A$  tracks the system's willingness to change  $A$  when you push it—the susceptibility. That link (fluctuation  $\leftrightarrow$  response) is what powers both Einstein's diffusion formula and the Green-Kubo relations.

$$\langle (\delta A)^2 \rangle \propto k_B T \cdot \chi_A \quad (\text{static susceptibility})$$

$$D = \mu_d \cdot k_B T \quad (\text{Einstein relation for diffusion})$$

$$L_{\{ab\}} = L_{\{ba\}} \quad (\text{Onsager reciprocity})$$

Plain English: diffusion speed  $D$  is the product of 'how easily a thing moves when pushed' (mobility  $\mu_d$ ) and the ambient trade-strength ( $k_B T$ ). Onsager's symmetry says there's no one-way trade channel near equilibrium: the cross-effects match in pairs.

## Transport from gradients (linear response)

Picking the right gradient: write the multiplier that weights routes for the conserved thing you care about, then take its gradient. That's why heat uses  $\nabla(1/T)$ , particles use  $-\nabla(\mu/T)$ , and momentum uses strain-rate. The rest are off-diagonal gates  $L_{\{ab\}}$ —paired by Onsager.

In VMS we do not start from constitutive laws. Transport is simply **route-traffic responding to small imbalances in cost**. Those imbalances are gradients of the multipliers that weight routes in the counting: when the local cost map tilts, traffic drifts. Thermodynamic "forces" are just those **multiplier gradients**; fluxes are the resulting **net traffic**.

Near equilibrium the drift is linear in the small tilts, and microreversibility of the route weights forces the coupling matrix to be symmetric. That gives the kernel:

$$J_a = \sum_b L_{ab} \cdot X_b$$

$$L_{ab} = L_{ba} \quad (\text{Onsager symmetry from time-reversal of route weights})$$

Here  $X_b$  are the **right gradients** (the multipliers that actually weight routes), and  $L_{\{ab\}}$  are **dimensionless** couplings that come from route-fluctuation cross-correlations (no new dials; calibration sets any material constants in the Appendix).

Where the L's come from: the transport matrix elements are integrated route-fluctuation cross-correlations (Green-Kubo style). That's why noise levels fix response strengths—same counting, two views.

Common force-flux pairs written in the VMS multipliers:

$$X_T = \nabla(1/T)$$

$$X_N = -\nabla(\mu/T)$$

$$X_{\text{mom}} \propto \nabla v \quad (\text{shear/strain-rate tensor})$$

So the examples you actually use are the same facts stated VMS-first: **heat flux**  $q$  responds to  $\nabla(1/T)$ ; **particle flux**  $J_N$  responds to  $-\nabla(\mu/T)$ ; **momentum flux** responds to strain-rate. Cross-effects (thermoelectric, Soret/Peltier) are just **off-diagonal gates**  $L_{\{ab\}}$ —they come in matched pairs by Onsager, because the route weights are reversible in the near-equilibrium limit.

What we publish: **scalings and ratios first** (many nuisances divide out), then calibrated numbers where it helps. No constitutive laws are imported; the forms above are bookkeeping consequences of route counting and the single-scale lock  $S_0 = \hbar$  (with  $k_B$  as an acceptance lock).

### Phase behavior (onset & stability without new dials)

Introduce an order parameter  $m$  and expand the free energy near onset. Stability is about signs and minima; nothing magical is added, and no new scales sneak in.

$$F(m) = F_0 + a(T) m^2 + b m^4 + \dots \quad (b > 0, \quad a(T_c) = 0)$$

If  $a(T)$  crosses zero, the minimum moves off  $m=0$  and a new phase appears. Coexistence and critical behavior follow from the same expansion; the details live in the Math Appendix.

## Bridges to Electromagnetism and Mechanics/GR

- EM bridge: the same levers determine medium response  $n(\lambda, T, \rho)$  and how noise spectra relate to dissipation (FDT). This isn't a bolt-on model: it's the same counting and the same derivatives.
- Mechanics/GR bridge: work/energy bookkeeping is identical; clocks and transport in background profiles use the shared kernels.

### Validity & limits

Near-equilibrium (linear response) and smooth-profile changes. Far-from-equilibrium and turbulent transport live in branches. Material-specific constants are calibrated; ratios across nearby states are usually tighter and preferred.

### Calibration capsule (how we compare to data)

- Targets: EoS checkpoints (compressibility, heat capacity), transport (thermal conductivity  $\kappa$ , viscosity  $\eta$ , diffusion  $D$ ), and phase thresholds for benchmark systems.
- Procedure: constants stamp  $\rightarrow$  derive scalings/ratios  $\rightarrow$  compare to canonical datasets (instrument/system notes logged).
- Why ratio-first: many nuisances divide out, so bands on  $\kappa(T_2)/\kappa(T_1)$  or  $C_p(T_2)/C_p(T_1)$  are more telling than a single absolute number.

### Predictions & falsifiers (high-level)

- Publish EoS and transport **\*\*ratio bands\*\*** across controlled  $T, p$  ranges; post calibrated curves as secondary.
- Falsifier: sustained  $\geq 3\sigma$  mismatch on multiple EoS or transport ratios in the same regime after systematics control (reject or bound the model region).

### Reader map (where to go next)

- Math Walk-Through — compact derivations for EoS, FDT/Onsager kernels; smooth-limit bridges included.
- Pillar Math Appendix — readable equations, symbol legend on page one, validity limits for each kernel.
- Calibration — constants stamp and datasets; ratio-vs-calibrated inserts and acceptance gates.
- Verification & Falsification — pass/fail tables with DOIs; pre-registered tests where possible.
- Student Workbook — pedagogy and worked problems, same notation as the Appendix.
- VMS Laws (Pilot) — concise cards you can use immediately, with classical limit boxes.

### Bottom line

One ontology, one scale ( $S_0 = \hbar$ ), and  $k_B$  as an acceptance lock. Thermodynamics here is route-counting plus free-energy levers and linear response—no new dials. We publish ratios first, calibrate where needed, and keep falsifiers explicit.