A Fiber Bundle Foundation for Classical Thermodynamics

Ripunjay Dwivedi[1,*]

[1] *Department of Physics, Indian Institute of Technology Bombay, India*

(Dated: April 2026)

The thermodynamic decomposition of energy transfer into work and heat is conventionally postulated as part of the first law. We show that it is instead a theorem of Newtonian mechanics and differential geometry, once the spatial confinement of an N-particle system is formalised as a fiber bundle over a manifold of macroscopic control parameters. A smooth family of diffeomorphisms $\varphi_\alpha : \Omega_0 \to \Omega(\alpha)$ trivialises the bundle and induces a flat connection whose horizontal–vertical decomposition splits the total differential of the internal energy into thermodynamic work (boundary deformation at frozen internal coordinates) and heat (internal rearrangement at fixed boundary): $\mathrm{d}U = \delta W + \delta Q$. This microscopic first law is the chain rule on a product space; it requires no equations of motion and no statistical input. To pass from microscopic to macroscopic, we show that linearity, positivity, and normalisation force the projection from the total phase space to the base manifold to be integration against a probability measure. Four physical constraints—support on the energy shell, flow-invariance, absolute continuity, and normalisation—together with the ergodic hypothesis, select the microcanonical measure uniquely. The canonical and grand canonical ensembles are then derived by marginalisation over bath degrees of freedom, with no additional postulates. Fiber averaging yields the macroscopic first law $\mathrm{d}U = \delta\bar{W} + \delta\bar{Q}$, exploiting the fact that U is constant on the energy shell while the generalized forces fluctuate. Re-expressing the macroscopic heat form as $\delta\bar{Q} = \mathrm{d}\Omega/\Sigma$—where Ω is the cumulative phase volume and Σ the density of states—reveals $1/T = k_B\Sigma/\Omega$ as the integrating factor and $S = k_B \ln \Omega$ as the entropy, yielding the Gibbs fundamental relation $\mathrm{d}U = T\,\mathrm{d}S + \delta\bar{W}$ without postulation. The Clausius inequality and the second law follow upon supplementing this framework with the Stosszahlansatz. Three assumptions beyond Newton's laws are required, each stated explicitly: a regularity hypothesis on the bundle structure, ergodicity, and molecular chaos. The framework delineates which features of classical thermodynamics are theorems of geometry and Hamiltonian mechanics, which depend on the choice of trivialisation, and which require statistical hypotheses.

I. INTRODUCTION

The first law of thermodynamics decomposes the change in internal energy into heat and work:

$$\mathrm{d}U = \delta Q + \delta W. \tag{1}$$

In standard treatments [10, 18, 44], this decomposition is postulated, motivated by the impossibility of perpetual motion machines of the first kind. The inexactness of δQ and δW—indicated by δ rather than d—is accepted as an empirical fact: the heat absorbed and work done along a thermodynamic process depend on the path, not merely on the endpoints.

This paper shows that the work–heat decomposition is not an independent postulate but a theorem of Newtonian mechanics and differential geometry, once a single physical observation is formalised: the particles of a thermodynamic system are *confined* to a domain whose geometry depends on externally controllable parameters. Entropy and temperature then emerge as explicit functions of the phase volume, and the second law follows from the molecular chaos hypothesis.

The logical chain is:

(1) Newton's laws and the work–energy theorem (no additional input);
(2) the internal energy $U = K_{\mathrm{rel}} + \Phi$ as a forced decomposition of microscopic work (no additional input);
(3) the fiber bundle structure of confinement (one regularity assumption);
(4) the microscopic first law $\mathrm{d}U = \delta W + \delta Q$ as the horizontal–vertical decomposition of $\mathrm{d}U$ (no additional input beyond step 3);
(5) the macroscopic first law $\mathrm{d}U = \delta\bar{W} + \delta\bar{Q}$ from fiber averaging (one dynamical assumption: ergodicity);
(6) entropy $S = k_B \ln \Omega$ and the Gibbs relation $\mathrm{d}U = T\,\mathrm{d}S + \delta\bar{W}$ (no additional input beyond step 5);
(7) the second law $\mathrm{d}S \geq \delta Q/T$ (one additional assumption: molecular chaos).

Three assumptions beyond Newton's laws are required. Each is stated explicitly and named: the bundle regularity hypothesis (Assumption V.1), ergodicity (Assumption VI.3), and the Stosszahlansatz (Remark VII.9). The paper does not claim to derive thermodynamics from mechanics alone; it claims that the work–heat decomposition, conventionally an independent postulate, is derivable once the physics of confinement is formalised geometrically.

The Central Construction

A gas is confined in a cylinder; its particles must lie within the cylinder walls. When the piston moves, the confining domain changes shape. When the piston is

* ripunjay.cmi@gmail.com

held fixed, the particles move freely within the domain. These are two geometrically distinct kinds of motion: one driven by the boundary, the other occurring within a fixed boundary.

The mathematical realisation of this distinction is a fiber bundle. The control parameters $\alpha = (\alpha_1, \ldots, \alpha_m)$: volume, length, plate separation, applied field—take values in a smooth manifold $\mathcal{B}$, the base. For each $\alpha \in \mathcal{B}$, the accessible N-particle configurations form a space $\mathcal{Q}_\alpha$, the fiber at α. The total configuration space $\mathcal{E}_\mathcal{Q} = \bigsqcup_\alpha \{\alpha\} \times \mathcal{Q}_\alpha$ is the total space, and the projection $\pi : \mathcal{E}_\mathcal{Q} \to \mathcal{B}$ records which control state each microscopic configuration belongs to. A smooth family of diffeomorphisms $\varphi_\alpha : \Omega_0 \to \Omega(\alpha)$—a *trivialization* of the bundle—then induces a product decomposition of the tangent space into base (horizontal) and fiber (vertical) directions, providing the connection through which the work–heat split is defined.

Applying this decomposition to the total differential of the internal energy $U = H_\alpha(\boldsymbol{\xi}, \boldsymbol{\pi})$—expressed in the canonical fiber coordinates $(\boldsymbol{\xi}, \boldsymbol{\pi})$ induced by the trivialization—yields the microscopic first law:

$$\mathrm{d}U = \underbrace{\sum_j \frac{\partial H_\alpha}{\partial \alpha_j}\bigg|_{\boldsymbol{\xi},\boldsymbol{\pi}} \mathrm{d}\alpha_j}_{\delta W \text{ (work)}} + \underbrace{\sum_i \frac{\partial H_\alpha}{\partial \boldsymbol{\xi}_i} \cdot \mathrm{d}\boldsymbol{\xi}_i + \sum_i \frac{\partial H_\alpha}{\partial \boldsymbol{\pi}_i} \cdot \mathrm{d}\boldsymbol{\pi}_i}_{\delta Q \text{ (heat)}}$$

This is the chain rule applied to a smooth function on a product space. The partition is sharp, coordinate-free (given the trivialization), and requires no equations of motion.

Relation to Existing Work

The geometric and symplectic machinery of classical mechanics [1, 4, 39] has been imported into thermodynamics through several channels. Hermann [25] and Mrugała [40, 41] identified the contact-geometric structure of the thermodynamic phase space. Weinhold [55] and Ruppeiner [49] introduced Riemannian metrics on the space of equilibrium states. Balian and Valentin [5] developed a fiber-bundle framework for thermodynamic potentials. Gay-Balmaz and Yoshimura [20] embedded non-equilibrium thermodynamics within geometric variational mechanics. Bravetti [9] and van der Schaft and Maschke [53] have developed contact and port-Hamiltonian formulations of dissipative and thermodynamic systems.

More recently, fiber bundles and gauge connections have entered the thermodynamics literature directly. Roberts [46] showed that heat defines a gauge connection on a line bundle over work configurations, with vanishing curvature equivalent to the local existence of entropy satisfying $\delta Q = T\,\mathrm{d}S$. Céleri and Rudnicki [12] introduced a gauge-invariant formulation of quantum thermodynamics, extended by Ferrari, Rudnicki, and Céleri [19] into a gauge theory with gauge-invariant entropy production.

The present paper differs from all of these in two respects. First, the bundle is not imposed on an already-constructed space of thermodynamic states; it *is* the mechanical configuration space of confined particles, arising from Newton's laws and the physics of confinement. The work–heat decomposition is derived rather than geometrised after the fact. Second, the treatment is entirely classical: no quantum mechanics, no information theory, and no prior thermodynamic definitions are used. The quantum-mechanical analogue of the construction—the Berry connection [6, 50] on the bundle of Hilbert spaces $\{L^2(\Omega(\alpha))\}_\alpha$—is a direct quantisation of the classical Ehresmann connection constructed here; the classical framework thus provides a substrate for the gauge-theoretic approaches cited above.

The axiomatic tradition in thermodynamics—Carathéodory [11], Lieb and Yngvason [36]—establishes the *existence* of entropy from topological or order-theoretic axioms, without reference to any microscopic model. The present approach is complementary: it does not prove existence abstractly but *exhibits* entropy as $S = k_B \ln \Omega$ and temperature as $1/T = k_B \Sigma/\Omega$, where Ω is the cumulative phase volume and Σ the density of states. The cost is model-dependence (the derivation relies on Hamiltonian mechanics); the gain is that S and T are computed as specific functions, not merely shown to exist.

Scope and Logical Status

The logical status of each ingredient should be stated precisely, to forestall misreading.

The *microscopic first law* (I) (Theorem V.7) is a theorem of classical mechanics and differential geometry. It requires Newton's laws, the fiber bundle structure of confinement (Assumption V.1), and the construction of the internal energy U as a state function (Definition IV.2). No statistical or thermodynamic input is needed.

The *macroscopic first law* $\mathrm{d}U = \delta\bar{W} + \delta\bar{Q}$ (Proposition VI.13) requires fiber averaging over an equilibrium probability measure. The projection from the total phase space $\mathcal{E}$ to the base $\mathcal{B}$ is forced by linearity, positivity, and normalisation to be integration against a probability measure (Proposition VI.1). The specific measure—the microcanonical—is selected by four physical constraints (support on the energy shell, flow-invariance, absolute continuity, normalisation) together with the ergodic hypothesis (Assumption VI.3).

The *entropy* $S = k_B \ln \Omega$ and the *Gibbs fundamental relation* $\mathrm{d}U = T\,\mathrm{d}S + \delta\bar{W}$ (Theorem VII.5) follow from re-expressing the macroscopic heat form $\delta\bar{Q}$ in terms of the phase volume Ω and the density of states Σ. No additional assumption is needed beyond those already used for the macroscopic first law.

The *second law* $\mathrm{d}S \geq \delta Q/T$ (Theorem VII.10) requires the molecular chaos assumption (Stosszahlansatz), which is the sole irreversible input in the entire derivation.

Organisation

Section II derives conservative forces from the demand that a conserved scalar exist, using a velocity-expansion argument that forces the conserved quantity to have the form $E = K + V$. Section III passes to N-particle systems, establishes König's theorem and the internal work–energy theorem. Section IV constructs the internal energy $U = K_{\rm rel} + \Phi$ by absorbing all conservative work into a potential. Section V is the geometric core: it constructs the fiber bundle, the trivialization, the canonical fiber coordinates, and the microscopic work–heat decomposition. Section VI introduces the projection problem, derives the microcanonical, canonical, and grand canonical measures from first principles, defines the fiber average, and proves the macroscopic first law. Section VII derives entropy and temperature from the phase volume, establishes the Gibbs fundamental relation and the Clausius inequality, and arrives at the second law.

a. Notation. Bold symbols denote vectors in $\mathbb{R}^3$. Indices j, k run over control parameters; Latin indices i run over particles. The symbol d denotes an exact differential; δ denotes an inexact 1-form—one that is not the differential of any smooth function on the relevant space. The summation convention is not used.

II. FROM NEWTON'S LAW TO CONSERVATIVE FORCES

The goal of this section is to understand, starting from Newton's Second Law alone, which force fields admit a conserved scalar quantity. We proceed by successively relaxing our demands until the algebra forces the concept of a conservative force upon us. The key step is an elementary argument (Lemma II.1) in which a putative first integral is expanded in powers of velocity and matched coefficient by coefficient against the Newtonian vector field; the resulting constraints kill every velocity-dependent term and force the conserved quantity to have the form $E = K + V$. We reproduce the argument in full because the identical logical pattern—a conservation demand forcing a gradient condition on the force—will recur in Section IV, where it leads to thermodynamics rather than to mechanics.

Momentum conservation: too rigid

Newton's Second Law for a particle of mass m in $\mathbb{R}^3$ reads

$$\frac{\mathrm{d}\mathbf{p}}{\mathrm{d}t} = \mathbf{F}(\mathbf{r}), \qquad \mathbf{p} \coloneqq m\dot{\mathbf{r}}. \tag{2}$$

The most immediate conserved quantity is the full momentum vector: $\mathbf{F} = \mathbf{0}$ implies $\mathbf{p} = \text{const}$. But this freezes both magnitude and direction of $\mathbf{p}$, excluding every non-trivial force. We ask for something weaker.

Conserving the magnitude: the work–energy theorem

Can $|\mathbf{p}|$ remain constant under a non-zero force? Since $\frac{\mathrm{d}}{\mathrm{d}t}|\mathbf{p}|^2 = 2m\,\dot{\mathbf{r}}\cdot\mathbf{F}$, the answer is: only if $\mathbf{F} \perp \dot{\mathbf{r}}$ everywhere: a genuine generalisation of $\mathbf{F} = \mathbf{0}$, but one that forbids the force from doing work. We want to track how $|\mathbf{p}|$ changes when work *is* done. Defining $K \coloneqq |\mathbf{p}|^2/(2m)$ and passing to differential form:

$$\mathrm{d}K = \mathbf{F}\cdot\mathrm{d}\mathbf{r} =: \delta W. \tag{3}$$

This is the **work–energy theorem**. The notation δW signals that the right-hand side need not be exact; it is a 1-form that may or may not be the differential of a function. The use of δ (or đ) to distinguish inexact from exact differentials dates to the 19th-century thermodynamic literature; for a modern treatment see [4, 22].

The apparent tension—an exact $\mathrm{d}K$ equalling a potentially inexact δW—is resolved by noting that the two sides live on different spaces. The kinetic energy $K(\mathbf{v})$ is a function on the velocity sector of the tangent bundle $\mathcal{P} = T\mathbb{R}^3$, while $\mathbf{F}\cdot\mathrm{d}\mathbf{r}$ is a 1-form on the position sector, pulled back to $\mathcal{P}$ via $(\mathbf{r}, \mathbf{v}) \mapsto \mathbf{r}$. The identity holds along each Newtonian trajectory—a one-dimensional submanifold of $\mathcal{P}$ on which every 1-form is trivially a total derivative [1, 4]. A closed loop in velocity space ($\oint \mathrm{d}K = 0$) need not correspond to a closed loop in position space, and vice versa, so there is no conflict.

The work–energy theorem is therefore an innocent on-shell identity: descriptive, not contradictory. A genuine contradiction can only come from imposing a further demand.

Conserving a general scalar: the gradient condition

We now impose the strongest demand compatible with non-trivial forces: the existence of a smooth function $E : \mathcal{P} \to \mathbb{R}$ satisfying

$$\dot{E} = 0 \quad \text{along every trajectory of (2).} \tag{4}$$

No assumption is made about the form of E.

Let $X_N = \mathbf{v}\cdot\nabla_{\mathbf{r}} + (\mathbf{F}/m)\cdot\nabla_{\mathbf{v}}$ be the Newtonian vector field on $\mathcal{P}$. By the existence and uniqueness theorem for ODEs [14], through every point $(\mathbf{r}_0, \mathbf{v}_0) \in \mathcal{P}$ there passes exactly one such trajectory, so the condition $X_N[E] = 0$ holds *at every point of* $\mathcal{P}$, not just along a particular orbit. Subtracting $X_N[K] = \mathbf{F}\cdot\mathbf{v}$:

$$X_N[\Phi] = -\,\mathbf{F}(\mathbf{r})\cdot\mathbf{v}, \qquad \Phi \coloneqq E - K, \tag{5}$$

for all $(\mathbf{r}, \mathbf{v}) \in \mathcal{P}$. The right-hand side is linear in $\mathbf{v}$. Since the equation must hold for *every* velocity at each position, it places severe constraints on the velocity dependence of Φ. We now show that these constraints kill every Taylor coefficient of Φ in $\mathbf{v}$ except the zeroth.

We call a force field *non-degenerate* if $\mathbf{F}(\mathbf{r})$ does not lie in a fixed proper subspace of $\mathbb{R}^3$ for all $\mathbf{r}$—every physically interesting three-dimensional force satisfies this.

The normalisation condition (6) below demands that the quadratic-in-$\mathbf{v}$ part of E is exactly the standard kinetic energy; relaxing it merely redefines an effective mass tensor and yields no new conservation law [28].

Lemma II.1. *Let* $\mathbf{F}$ *be smooth and non-degenerate. Suppose* $\Phi = E - K$ *satisfies* (5) *and the kinetic normalisation*

$$\left.\frac{\partial^2 E}{\partial v_i\,\partial v_j}\right|_{\mathbf{v}=0} = m\,\delta_{ij}. \tag{6}$$

Then Φ *depends only on position:* $\Phi(\mathbf{r},\mathbf{v}) = \Phi_0(\mathbf{r})$.

Proof. Expand $\Phi = \sum_{n\geq 0}\Phi_n$ in homogeneous polynomials of degree n in $\mathbf{v}$. Substituting into (5), the left-hand side becomes a power series in $\mathbf{v}$; the right-hand side is purely linear. Matching degree-d monomials gives, for each d,

$$\sum_k v_k\,\partial_{r_k}\Phi_{d-1} \quad + \quad \frac{d+1}{m}\sum_k F_k\,\Phi_{d+1}^{k\,i_1\cdots i_d}\,v_{i_1}\cdots v_{i_d} = \begin{cases} -\sum_i F_i\,v_i & d=1, \\ 0 & d\neq 1, \end{cases} \tag{7}$$

where $\Phi_{d+1}^{k\,i_1\cdots i_d}$ denotes the fully symmetric coefficient tensor of the degree-$(d+1)$ homogeneous polynomial Φ_{d+1} in $\mathbf{v}$, contracted with F_k over the index k.
Step 1: At $d=0$: $(\mathbf{F}/m)\cdot\boldsymbol{\Phi}_1 = 0$ everywhere; non-degeneracy forces $\Phi_1 = 0$.
Step 2: The normalisation (6) gives $m\delta_{ij} + 2\Phi_{2,ij} = m\delta_{ij}$, so $\Phi_2 = 0$.
Step 3: Two interleaved inductions kill all higher coefficients. The key observation: if $\Phi_{d-1} = 0$ for some $d \geq 2$, the first term in (7) vanishes, and the second gives $(d+1)(F_k/m)\,\Phi_{d+1,k\cdots} = 0$ for all $\mathbf{r}$, forcing $\Phi_{d+1} = 0$ by non-degeneracy. Inducting from $\Phi_1 = 0$ via $d = 2,4,6,\ldots$ gives $\Phi_3 = \Phi_5 = \Phi_7 = \cdots = 0$; from $\Phi_2 = 0$ via $d = 3,5,7,\ldots$ gives $\Phi_4 = \Phi_6 = \Phi_8 = \cdots = 0$. □

The upshot: insisting on a conserved energy forces $\Phi = E - K$ to depend on position alone. Substituting $\Phi = \Phi_0(\mathbf{r})$ back into (5):

$$\nabla\Phi_0\cdot\mathbf{v} = -\,\mathbf{F}\cdot\mathbf{v} \quad \text{for all } \mathbf{v}, \tag{8}$$

whence $\nabla\Phi_0 = -\mathbf{F}$ and therefore

$$\nabla\times\mathbf{F} = -\nabla\times(\nabla\Phi_0) = \mathbf{0}. \tag{9}$$

This is where the contradiction lives. The conservation demand has forced a curl condition on $\mathbf{F}$; any force with $\nabla\times\mathbf{F}\neq\mathbf{0}$ at even a single point is *provably incompatible* with the existence of a conserved E, regardless of the functional form allowed. The work–energy theorem played no role in producing this contradiction—it was an innocent bystander throughout; the real engine was the demand (4).

The converse is classical:

Theorem II.2 (Poincaré Lemma [45]; see [8, 54])**.** *A smooth* 1*-form* $\mathbf{F}\cdot d\mathbf{r}$ *on a simply connected open* $U\subseteq\mathbb{R}^3$ *is exact if and only if* $\nabla\times\mathbf{F}=\mathbf{0}$ *on* U.

Definition II.3 (Conservative Force and Potential Energy)**.** A smooth force $\mathbf{F}: U\to\mathbb{R}^3$ on a simply connected open set U is *conservative* if $\nabla\times\mathbf{F}=\mathbf{0}$. The Poincaré Lemma guarantees a unique (up to a constant) function $V: U\to\mathbb{R}$ with $\mathbf{F} = -\nabla V$.

Identifying $\Phi_0 = V$, the conserved quantity is forced to be $E = K + V$—it could not take any other form.

Theorem II.4 (Conservation of Mechanical Energy)**.** *If* $\mathbf{F} = -\nabla V$ *on a simply connected* U*, then* $E = K + V$ *satisfies* $\dot{E} = 0$ *along every trajectory of* (2).

Proof. $\dot{E} = \dot{K} + \dot{V} = \mathbf{F}\cdot\dot{\mathbf{r}} - \mathbf{F}\cdot\dot{\mathbf{r}} = 0$. □

Remark II.5 (Simply connected topology). On non-simply-connected domains, $\nabla\times\mathbf{F}=\mathbf{0}$ remains *necessary* for conservation; *sufficiency* fails because the global potential may be multi-valued. This underlies the theory of magnetic monopoles [15, 56] and the Aharonov–Bohm effect [2]; for the fiber bundle perspective, see [42].

Remark II.6 (Relation to prior work on polynomial integrals). The general programme of classifying first integrals of natural mechanical systems that are polynomial in the velocities has a substantial literature; see Kozlov [34] for a systematic treatment of integrability and non-integrability of Hamiltonian systems, and Hojman and Harleston [28] for the construction of conserved quantities under generalised kinetic normalisation. The specific argument of Lemma II.1—the two interleaved inductions exploiting the parity structure of the degree-matching equations—appears to be new in this form.

III. MULTI-PARTICLE SYSTEMS: KÖNIG'S THEOREM AND THE INTERNAL WORK–ENERGY THEOREM

Section II established the work–energy structure for a single particle. We now pass to N interacting particles. The new question is whether the kinetic energy can be separated into a bulk-translation piece and an internal-motion piece, each obeying its own energy balance. This separation is the kinematic foundation on which thermodynamics rests: the distinction between ordered macroscopic kinetic energy and disordered thermal energy is meaningful only if the two can be cleanly disentangled.

König's Theorem

Consider N particles with masses $m_i > 0$ and positions $\mathbf{r}_i(t)$. Define the center-of-mass (COM) quantities:

$$\mathbf{R} \coloneqq \frac{1}{M}\sum_i m_i\mathbf{r}_i, \quad \mathbf{V}\coloneqq\dot{\mathbf{R}}, \quad M\coloneqq\sum_i m_i. \tag{10}$$

Decomposing $\mathbf{r}_i = \mathbf{R} + \mathbf{r}'_i$ and $\mathbf{v}_i = \mathbf{V} + \mathbf{v}'_i$, the COM constraints $\sum_i m_i \mathbf{r}'_i = \mathbf{0}$ and $\sum_i m_i \mathbf{v}'_i = \mathbf{0}$ hold identically. Expanding $K_{\text{total}} = \sum_i \frac{1}{2} m_i |\mathbf{V} + \mathbf{v}'_i|^2$, the cross term $\mathbf{V} \cdot \sum_i m_i \mathbf{v}'_i$ vanishes by the COM constraint:

Theorem III.1 (König's Theorem [33]; see [22])**.**

$$K_{\text{total}} = \underbrace{\tfrac{1}{2} M |\mathbf{V}|^2}_{K_{\text{macro}}} + \underbrace{\sum_i \tfrac{1}{2} m_i |\mathbf{v}'_i|^2}_{K_{\text{rel}}} . \tag{11}$$

The separation is exact, and not an approximation. Geometrically, it is the Pythagorean theorem for the decomposition of the system velocity in $\mathbb{R}^{3N}$ into a 3-dimensional uniform-translation component and a $(3N-3)$-dimensional internal component, orthogonal with respect to the mass-weighted inner product

$$\langle \mathbf{u}, \mathbf{w} \rangle_m = \sum_i m_i \mathbf{u}_i \cdot \mathbf{w}_i .$$

The internal work–energy theorem

Newton's equations $m_i \ddot{\mathbf{r}}_i = \mathbf{F}_i$ give the total work–energy theorem $\mathrm{d}K_{\text{total}} = \sum_i \mathbf{F}_i \cdot \mathrm{d}\mathbf{r}_i$. Summing the equations of motion and using Newton's Third Law (internal forces cancel because they multiply the *same* displacement $\mathrm{d}\mathbf{R}$):

$$M\ddot{\mathbf{R}} = \sum_i \mathbf{F}_i^{\text{ext}}, \qquad \mathrm{d}K_{\text{macro}} = \Big(\sum_i \mathbf{F}_i\Big) \cdot \mathrm{d}\mathbf{R}. \tag{12}$$

Writing $\mathrm{d}\mathbf{r}_i = \mathrm{d}\mathbf{R} + \mathrm{d}\mathbf{r}'_i$, the total work splits into a macroscopic part $(\sum_i \mathbf{F}_i) \cdot \mathrm{d}\mathbf{R}$ and a microscopic part $\sum_i \mathbf{F}_i \cdot \mathrm{d}\mathbf{r}'_i$. Since the König constraint holds at every instant, differentiating (11) gives $\mathrm{d}K_{\text{total}} = \mathrm{d}K_{\text{macro}} + \mathrm{d}K_{\text{rel}}$. Subtracting the macroscopic theorem:

Proposition III.2 (Internal Work–Energy Theorem)**.**

$$\mathrm{d}K_{\text{rel}} = \sum_i \mathbf{F}_i \cdot \mathrm{d}\mathbf{r}'_i =: \delta W_{\text{micro}}. \tag{13}$$

The internal sector has its own self-contained energy balance, completely decoupled from the bulk translation. As with (3), the notation $\mathrm{d}K_{\text{rel}} = \delta W_{\text{micro}}$ is an on-shell identity and raises no contradiction; a genuine contradiction can only come from imposing an additional conservation demand, which is the subject of the next section.

Note that internal forces do *not* cancel in δW_{micro}: the pair $(\mathbf{F}_{ij}, \mathbf{F}_{ji} = -\mathbf{F}_{ij})$ now multiplies *different* relative displacements $(\mathrm{d}\mathbf{r}'_i, \mathrm{d}\mathbf{r}'_j)$, so the cancellation that occurred in the COM equation fails. This is essential: inter-particle forces are the mechanism by which energy moves between the macroscopic and microscopic sectors. If they cancelled, the internal energy could never change and thermodynamics would be trivial.

IV. THE SECOND EXACTNESS CRISIS: THERMODYNAMIC INTERNAL ENERGY

The internal work–energy theorem has the same structural form as the single-particle case: exact differential on the left, potentially inexact 1-form on the right. As before, this is not a contradiction. A contradiction arises only when we impose a conservation demand—and the chain it triggers will now run into a wall that forces us to *weaken* the demand, discovering the thermodynamic internal energy in the process.

Demanding full conservation—and the contradiction it produces

Consider the N-particle system of Section III, and let $\mathcal{P}_{\text{rel}} \coloneqq T\mathcal{Q}_{\text{rel}}$ denote the phase space of the internal degrees of freedom, where

$$\mathcal{Q}_{\text{rel}} = \{(\mathbf{r}'_1, \ldots, \mathbf{r}'_N) : \sum_i m_i \mathbf{r}'_i = \mathbf{0}\}$$

is the relative configuration space. Does there exist a smooth function $U : \mathcal{P}_{\text{rel}} \to \mathbb{R}$ with $\dot{U} = 0$ along every internal trajectory? We assume nothing about its form.

The argument of Section II carries over to the $(3N-3)$-dimensional internal phase space. Non-degeneracy generalises naturally:

Definition IV.1 (Non-degeneracy: N-particle internal forces)**.** The force system $\{\mathbf{F}_i\}$ is *non-degenerate on* $\mathcal{Q}_{\text{rel}}$ if the projected forces span the $(3N-3)$-dimensional relative-force subspace $\{\sum_i m_i \mathbf{f}_i = \mathbf{0}\}$ as $\mathbf{r}'$ varies.

Expanding $\Psi \coloneqq U - K_{\text{rel}}$ in the relative velocities and matching order by order against the identity $X_{\text{int}}[\Psi] = -\sum_i \mathbf{F}_i \cdot \mathbf{v}'_i$ forces $\Psi = \Psi(\mathbf{r}')$ by the same two-induction argument as Lemma II.1. Matching the linear-in-$\mathbf{v}'$ coefficients then gives

$$\nabla_{\mathbf{r}'_i} \Psi = -\mathbf{F}_i \quad \text{for each } i \qquad \Longrightarrow \frac{\partial F_{i,a}}{\partial r'_{j,b}} = \frac{\partial F_{j,b}}{\partial r'_{i,a}} \quad \forall\, i, j, a, b. \tag{14}$$

This is the same logical pattern as Section II: the conservation demand has forced a curl-free condition on every force.

In the single-particle case, this was the end of the story. Here, it is a dead end. Real thermodynamic systems—gases, liquids, solids—include velocity-dependent dissipation, time-varying fields, and radiation reaction: forces that violate (14). For such systems, no conserved U exists. Insisting on full conservation would restrict us to systems where every internal force is conservative, excluding virtually everything of thermodynamic interest.

The resolution is not to abandon the search for an energy function, but to ask a weaker question: *can we absorb as much of δW_{micro} as possible into the differential of a state function, and cleanly isolate what remains?*

Absorbing the exact part: the forced decomposition

The curl-free condition (14) draws a sharp, forced line via the Poincaré Lemma (Theorem II.2):

(i) **Microscopically conservative forces** satisfy (14) and admit a potential: $\mathbf{F}_i^{\text{cons}} = -\nabla_{\mathbf{r}'_i}\Phi$ for some $\Phi(\mathbf{r}')$. This class includes Coulomb, Lennard-Jones, and gravitational interactions, and static external potentials. Their work is exact: $\sum_i \mathbf{F}_i^{\text{cons}} \cdot \mathrm{d}\mathbf{r}'_i = -\mathrm{d}\Phi$.

(ii) **Non-conservative forces** do not satisfy (14): velocity-dependent dissipation, radiation reaction, etc. No potential exists; their work is genuinely inexact.

This is not a modelling choice—it is the unique decomposition forced by the Poincaré Lemma. Decomposing the internal work accordingly and substituting into the internal work–energy theorem:

$$\begin{aligned} \mathrm{d}K_{\text{rel}} &= -\mathrm{d}\Phi + \delta W_{\text{nc}} \\ \Longrightarrow \quad \mathrm{d}(K_{\text{rel}} + \Phi) &= \delta W_{\text{nc}}. \end{aligned} \tag{15}$$

Definition IV.2 (Thermodynamic Internal Energy)**.**

$$U \coloneqq K_{\text{rel}} + \Phi(\mathbf{r}'), \qquad \mathrm{d}U = \delta W_{\text{nc}}. \tag{16}$$

U is a well-defined state function on $\mathcal{P}_{\text{rel}}$: its value depends only on the current microstate, not on history. It is *not* conserved when non-conservative forces are present—$\mathrm{d}U = \delta W_{\text{nc}} \neq 0$ in general. If all forces are conservative, $\delta W_{\text{nc}} = 0$ and full conservation is recovered as a special case.

The equation $\mathrm{d}U = \delta W_{\text{nc}}$ is not a contradiction. U is an explicitly constructed smooth function; its differential is automatically exact by the tautology of calculus. There is nothing paradoxical in a smooth function having a non-zero rate of change—that is simply what "U is not conserved" means. The contradiction of Section IV arose from demanding $\dot{U} = 0$; dropping that demand removes the contradiction entirely, and what remains is the cleanest energy balance the forces allow.

Remark IV.3 (The unifying thread). The same logical pattern has now appeared three times: $\mathrm{d}K = \delta W$ in Section II, $\mathrm{d}K_{\text{rel}} = \delta W_{\text{micro}}$ in Section III, and $\mathrm{d}U = \delta W_{\text{nc}}$ here. In each case, writing "exact = potentially inexact" is never by itself a contradiction—the tension is notational, not logical. A genuine contradiction arises only when a conservation demand is imposed and the resulting gradient condition conflicts with the actual forces. When the demand is relaxed, the forced decomposition $U = K_{\text{rel}} + \Phi$ with $\mathrm{d}U = \delta W_{\text{nc}}$ is what remains.

V. THE GEOMETRY OF CONFINEMENT AND THE WORK–HEAT DECOMPOSITION

The internal configuration space $\mathcal{Q}_{\text{rel}}$ has dimension $3N-3$, far too large for macroscopic description. In practice, a thermodynamic system is characterised by a small number of *macroscopic observables*—volume, length, plate separation, magnetisation—each of which is determined by the microscopic configuration $(\mathbf{r}'_1, \ldots, \mathbf{r}'_N)$. Some observables are smooth functions on $\mathcal{Q}_{\text{rel}}$: the end-to-end length of a polymer is $L = |\mathbf{r}'_N - \mathbf{r}'_1|$, and the magnetisation of a classical spin system is $M = \sum_i \mu_i(\mathbf{r}'_i)$. Others, such as the volume of a confining region, characterise the domain Ω to which the $\mathbf{r}'_i$ are restricted rather than being pointwise functions of the configuration.

We now make the following assumption, which is the additional structure we impose to pass from mechanics to thermodynamics.

Assumption V.1 (Macroscopic control parameters)**.** There exist m independent quantities $\alpha_1, \ldots, \alpha_m$, with $m \ll 3N-3$, such that:

(a) each α_j is experimentally accessible and externally controllable;
(b) the accessible region of $\mathcal{Q}_{\text{rel}}$ at fixed α is a well-defined subset $\mathcal{Q}_\alpha \subset \mathcal{Q}_{\text{rel}}$;
(c) the family $\{\mathcal{Q}_\alpha\}_{\alpha\in\mathcal{B}}$, where $\mathcal{B} \subseteq \mathbb{R}^m$ is the space of attainable control states, varies smoothly with α in the sense that the total space

$$\mathcal{E}_{\mathcal{Q}} \coloneqq \bigsqcup_{\alpha\in\mathcal{B}} \{\alpha\} \times \mathcal{Q}_\alpha \tag{17}$$

admits the structure of a smooth fiber bundle under the projection $\pi_{\mathcal{Q}} : \mathcal{E}_{\mathcal{Q}} \to \mathcal{B}$.

Conditions (a) and (b) are physical: they encode the empirical fact that a macroscopic observer can fix the volume of a container, the length of a wire, or the separation of capacitor plates, and that doing so determines which microscopic configurations are accessible. Condition (c) is a regularity hypothesis: it asserts that the dependence of the accessible configuration space on the control parameters is smooth enough to support differential-geometric constructions. This condition is satisfied by all standard thermodynamic systems—uniformly scaled boxes, isotropically compressed spheres, linearly stretched wires—and fails only for pathological boundary geometries (fractal containers, topological transitions) that lie outside the scope of classical equilibrium thermodynamics. Condition (c) requires the confining domains $\Omega(\alpha)$ to be mutually diffeomorphic for all $\alpha \in \mathcal{B}$ and $\partial\Omega(\alpha)$ to vary smoothly as a family of embedded submanifolds of $\mathbb{R}^3$; the N-particle fibers $\mathcal{Q}_\alpha \subset \mathbb{R}^{3N-3}$ then inherit smooth variation in α.

The idea that geometric structures on the space of thermodynamic states can illuminate the foundations of the subject has a substantial history. Hermann [25] introduced contact geometry into equilibrium thermodynamics; Mrugała [40, 41] developed the contact-geometric formulation systematically; and Balian and Valentin [5] constructed a fiber-bundle framework for thermodynamic potentials. More recently, Gay-Balmaz and Yoshimura [20] have embedded non-equilibrium thermodynamics within the variational structure of geometric mechanics. The contact-geometric approach to dissipative systems has been extended to thermodynamic settings by Bravetti and collaborators [9], and van der Schaft and Maschke [53] have developed a port-Hamiltonian formulation of thermodynamic systems using Dirac structures.

The present treatment differs from these in starting from the N-particle configuration space rather than from a pre-existing thermodynamic state space: the fiber bundle here is the *mechanical* configuration space of confined particles, not an abstract manifold of macroscopic variables.

The Fibers: Accessible Configurations at Fixed Control State

Let $\{\alpha_j\}_{j=1}^m$ denote the control parameters of Assumption V.1, taking values in a smooth m-dimensional manifold $\mathcal{B}$—the *base manifold*. Representative examples include: volume V (gas in cylinder), length L (stretched wire), plate separation d (capacitor), strain tensor ϵ_{jk} (elastic solid), and applied field H (magnetic system).

For each $\alpha \in \mathcal{B}$, the particles are confined to a domain $\Omega(\alpha) \subset \mathbb{R}^3$, and the accessible N-particle relative configuration space is

$$\mathcal{Q}_\alpha \coloneqq \Big\{ (\mathbf{r}'_1, \ldots, \mathbf{r}'_N) \in (\mathbb{R}^3)^N \Big| \sum_i m_i \mathbf{r}'_i = \mathbf{0}, \quad \mathbf{r}'_i \in \Omega(\alpha) \quad \forall i \Big\}. \tag{18}$$

The *phase-space fiber* at α is $\mathcal{F}_\alpha \coloneqq T^*\mathcal{Q}_\alpha$, with coordinates

$$(\mathbf{r}', \mathbf{p}') \coloneqq (\mathbf{r}'_1, \ldots, \mathbf{r}'_N, \mathbf{p}'_1, \ldots, \mathbf{p}'_N),$$

where $\mathbf{p}'_i = m_i \mathbf{v}'_i$. We carefully distinguish the *configuration fiber* $\mathcal{Q}_\alpha$ from the *phase-space fiber* $\mathcal{F}_\alpha$ throughout.

The total configuration space and total phase space are

$$\mathcal{E}_\mathcal{Q} \coloneqq \bigsqcup_{\alpha \in \mathcal{B}} \{\alpha\} \times \mathcal{Q}_\alpha, \qquad \mathcal{E} \coloneqq \bigsqcup_{\alpha \in \mathcal{B}} \{\alpha\} \times \mathcal{F}_\alpha, \tag{19}$$

with projections $\pi_\mathcal{Q} : \mathcal{E}_\mathcal{Q} \to \mathcal{B}$ and $\pi : \mathcal{E} \to \mathcal{B}$.

The internal energy $U = K_{\text{rel}} + \Phi(\mathbf{r}')$ of Definition IV.2 is now a smooth function on the total phase space $\mathcal{E}$. Its value at a point $(\alpha, \mathbf{r}', \mathbf{p}') \in \mathcal{E}$ depends on the full microstate: two configurations in the same fiber—same α, different particle arrangements—generically have different kinetic energies and different values of the inter-particle potential. In particular, U does not factor through the projection π: there is no function $\bar{U} : \mathcal{B} \to \mathbb{R}$ with $U = \bar{U} \circ \pi$.

To decompose $\mathrm{d}U$ into a part driven by changes in α and a part driven by internal rearrangement, one must be able to compare fibers at different values of α. But without further structure, this comparison is ill-defined: the fiber coordinate $\mathbf{r}'$ at α ranges over $\mathcal{Q}_\alpha \subset \Omega(\alpha)^N$, while at α' it ranges over $\mathcal{Q}_{\alpha'} \subset \Omega(\alpha')^N$—these are different subsets of $(\mathbb{R}^3)^N$, so the statement "$\mathbf{r}'$ is held fixed while α varies" is meaningless. A trivialization $\varphi_\alpha : \Omega_0 \to \Omega(\alpha)$ resolves this by introducing canonical coordinates $\boldsymbol{\xi}_i \in \Omega_0$ that are α-independent: after writing $\mathbf{r}'_i = \varphi_\alpha(\boldsymbol{\xi}_i)$, changes in α at frozen $\boldsymbol{\xi}$ are well-defined.

The Trivialization: Comparing Fibers Across Control States

Let Ω_0 denote a bounded reference domain. Since $\Omega(\alpha)$ is bounded for every α (the particles are confined), this entails no loss of generality. We assume the existence of a smooth family of diffeomorphisms

$$\varphi_\alpha : \Omega_0 \longrightarrow \Omega(\alpha), \qquad \alpha \in \mathcal{B}, \tag{20}$$

which identifies each point $\boldsymbol{\xi}_i \in \Omega_0$ with the physical position $\mathbf{r}'_i = \varphi_\alpha(\boldsymbol{\xi}_i) \in \Omega(\alpha)$. The *canonical internal coordinates* $\boldsymbol{\xi}_i \in \Omega_0$ are dimensionless and boundary-agnostic; all geometric information about the confining domain is encoded in φ_α.

The map (20) induces a global trivialization $\Phi_{\text{triv}} : \mathcal{B} \times \Omega_0^N \to \mathcal{E}_\mathcal{Q}$ by $\Phi_{\text{triv}}(\alpha, \boldsymbol{\xi}) = (\alpha, \varphi_\alpha(\boldsymbol{\xi}_1), \ldots, \varphi_\alpha(\boldsymbol{\xi}_N))$, which is a diffeomorphism by the smoothness of φ_α and the inverse function theorem. In particular, $(\mathcal{E}_\mathcal{Q}, \pi_\mathcal{Q}, \mathcal{B})$ is a smooth fiber bundle.

Remark V.2 (Affine scaling as a special case). For domains obtained by applying a linear deformation $A(\alpha) \in \mathrm{GL}(3, \mathbb{R})$ to the reference domain, $\varphi_\alpha(\boldsymbol{\xi}) = A(\alpha)\boldsymbol{\xi}$. For a cubical box of volume $V = L^3$, this reduces to $\varphi_\alpha(\boldsymbol{\xi}_i) = V^{1/3}\boldsymbol{\xi}_i$. The affine case admits explicit computation of the fiber metric and the thermodynamic forces, but none of the structural results that follow depend on it.

With the trivialization in hand, the total differential of $\mathbf{r}'_i = \varphi_\alpha(\boldsymbol{\xi}_i)$ separates by the chain rule into two geometrically distinct contributions:

$$\mathrm{d}\mathbf{r}'_i = \underbrace{\sum_j \frac{\partial \varphi_\alpha(\boldsymbol{\xi}_i)}{\partial \alpha_j} \, \mathrm{d}\alpha_j}_{\mathrm{d}\mathbf{r}'^{\,H}_i \text{ (horizontal)}} + \underbrace{\frac{\partial \varphi_\alpha(\boldsymbol{\xi}_i)}{\partial \boldsymbol{\xi}_i} \cdot \mathrm{d}\boldsymbol{\xi}_i}_{\mathrm{d}\mathbf{r}'^{\,V}_i \text{ (vertical)}}. \tag{21}$$

The horizontal component is the displacement induced by a change $\mathrm{d}\alpha_j$ in the control parameters at frozen internal coordinates $\boldsymbol{\xi}_i$: the boundary deforms and carries the particles passively with it. The vertical component is the displacement at the frozen boundary ($\mathrm{d}\alpha = 0$): the particles rearrange within the fixed domain. This is the **fundamental kinematic split**, and it is an instance of the Ehresmann horizontal–vertical decomposition on the bundle $(\mathcal{E}_\mathcal{Q}, \pi_\mathcal{Q}, \mathcal{B})$ induced by the trivialization.

Substituting (21) into $\mathrm{d}U$ now produces the sought decomposition. To do this cleanly, we first pass to canonical fiber coordinates on the full phase-space bundle $\mathcal{E}$.

The kinematic split (21) decomposes displacements in the configuration-space bundle $\mathcal{E}_\mathcal{Q}$. To decompose $\mathrm{d}U$—which depends on both positions and momenta—we must extend the trivialization to the phase-space bundle $\mathcal{E} = \bigsqcup_\alpha \{\alpha\} \times \mathcal{F}_\alpha$.

The physical momenta $\mathbf{p}'_i = m_i \dot{\mathbf{r}}'_i$ are not canonically conjugate to the internal coordinates $\boldsymbol{\xi}_i$; the correct conjugate momenta are obtained from the trivialization map.

Definition V.3 (Canonical fiber coordinates)**.** Let $J_\alpha(\boldsymbol{\xi}_i) \coloneqq \partial\varphi_\alpha(\boldsymbol{\xi}_i)/\partial\boldsymbol{\xi}_i$ denote the Jacobian of the trivialization map at $\boldsymbol{\xi}_i$. The *canonical fiber momentum* conjugate to $\boldsymbol{\xi}_i$ is

$$\boldsymbol{\pi}_i \;\coloneqq\; J_\alpha(\boldsymbol{\xi}_i)^T\,\mathbf{p}'_i. \tag{22}$$

The pair $(\boldsymbol{\xi}_i, \boldsymbol{\pi}_i)$ constitutes a Darboux chart on the fiber. The control parameters α_j are not dynamical variables and carry no conjugate momenta; the symplectic form $\omega_{\text{symp}} = \sum_i \mathrm{d}\boldsymbol{\xi}_i \wedge \mathrm{d}\boldsymbol{\pi}_i$ is defined on each fiber $\mathcal{F}_\alpha$, not on the total space $\mathcal{E}$.

The internal energy $U = K_{\text{rel}} + \Phi$, re-expressed in these coordinates, becomes the *fiber Hamiltonian*:

Definition V.4 (Fiber Hamiltonian)**.**

$$\begin{aligned} H_\alpha(\boldsymbol{\xi},\boldsymbol{\pi}) \;&\coloneqq\; K(\boldsymbol{\xi},\boldsymbol{\pi};\alpha) \;+\; \Phi_{\text{full}}(\boldsymbol{\xi};\alpha), \\ K(\boldsymbol{\xi},\boldsymbol{\pi};\alpha) \;&\coloneqq\; \sum_{i=1}^{N} \frac{\boldsymbol{\pi}_i^T\, g_\alpha(\boldsymbol{\xi}_i)^{-1}\,\boldsymbol{\pi}_i}{2m_i}, \\ \Phi_{\text{full}}(\boldsymbol{\xi};\alpha) \;&\coloneqq\; \Phi\big(\varphi_\alpha(\boldsymbol{\xi}_1),\ldots,\varphi_\alpha(\boldsymbol{\xi}_N)\big), \end{aligned} \tag{23}$$

where $g_\alpha(\boldsymbol{\xi}_i) \coloneqq J_\alpha(\boldsymbol{\xi}_i)^T J_\alpha(\boldsymbol{\xi}_i)$ is the *metric tensor* on the fiber induced by the trivialization.

Remark V.5 (Both sectors depend on α). In the physical coordinates $(\mathbf{r}', \mathbf{p}')$, the kinetic energy $K_{\text{rel}} = \sum_i |\mathbf{p}'_i|^2/(2m_i)$ is independent of the control parameters, and the potential $\Phi(\mathbf{r}')$ carries no intrinsic α-dependence. In the canonical coordinates $(\boldsymbol{\xi}, \boldsymbol{\pi})$, both acquire explicit α-dependence: the kinetic energy through the metric g_α, and the potential through the composition with φ_α. This α-dependence is not an artefact of the coordinates; it encodes the physical fact that boundary deformation changes both the kinetic energy of particles carried passively by the trivialization map and the energy landscape seen in the canonical coordinates.

Corollary V.6 (Flatness)**.** *The connection induced by the trivialization has vanishing curvature. This is immediate: in the product coordinates* $(\alpha, \boldsymbol{\xi}, \boldsymbol{\pi})$*, horizontal vector fields have no fiber components, so their Lie bracket has no fiber components, and the curvature* 2*-form vanishes identically.*

The Work–Heat Decomposition

The internal energy on the total space is $U = H_\alpha(\boldsymbol{\xi}, \boldsymbol{\pi})$—a smooth function of the coordinates $(\alpha, \boldsymbol{\xi}, \boldsymbol{\pi})$ on $\mathcal{E}$. Its total differential is, by the chain rule,

$$\begin{aligned} \mathrm{d}U = &\underbrace{\sum_{j=1}^{m} \frac{\partial H_\alpha}{\partial \alpha_j}\bigg|_{\boldsymbol{\xi},\boldsymbol{\pi}} \mathrm{d}\alpha_j}_{\text{base directions}} \\ &+ \underbrace{\sum_{i=1}^{N} \frac{\partial H_\alpha}{\partial \boldsymbol{\xi}_i} \cdot \mathrm{d}\boldsymbol{\xi}_i + \sum_{i=1}^{N} \frac{\partial H_\alpha}{\partial \boldsymbol{\pi}_i} \cdot \mathrm{d}\boldsymbol{\pi}_i}_{\text{fiber directions}}. \end{aligned} \tag{24}$$

In the trivialized coordinates $(\alpha, \boldsymbol{\xi}, \boldsymbol{\pi})$, the tangent space to $\mathcal{E}$ at every point splits naturally into *base directions* (spanned by $\partial/\partial\alpha_j$) and *fiber directions* (spanned by $\partial/\partial\boldsymbol{\xi}_i$ and $\partial/\partial\boldsymbol{\pi}_i$): this is simply the product structure $T(\mathcal{B}\times\mathcal{F}) = T\mathcal{B} \oplus T\mathcal{F}$ of the trivialized bundle. A tangent vector is called *horizontal* if its fiber components vanish ($\mathrm{d}\boldsymbol{\xi}_i = 0$, $\mathrm{d}\boldsymbol{\pi}_i = 0$ for all i) and *vertical* if its base components vanish ($\mathrm{d}\alpha_j = 0$ for all j). Equation (24) is then the decomposition of $\mathrm{d}U$ into its horizontal and vertical parts. (In the language of differential geometry, this horizontal–vertical splitting is the *Ehresmann connection* on $\mathcal{E}$ induced by the trivialization [17, 32].)

Theorem V.7 (Work–Heat Decomposition)**.**

$$\mathrm{d}U \;=\; \delta W + \delta Q, \tag{25}$$

where

$$\delta W \coloneqq \sum_{j=1}^{m} \frac{\partial H_\alpha}{\partial \alpha_j}\bigg|_{\boldsymbol{\xi},\boldsymbol{\pi}} \mathrm{d}\alpha_j \qquad \text{(horizontal: thermodynamic work)}, \tag{26}$$

$$\delta Q \coloneqq \sum_{i=1}^{N} \frac{\partial H_\alpha}{\partial \boldsymbol{\xi}_i} \cdot \mathrm{d}\boldsymbol{\xi}_i + \sum_{i=1}^{N} \frac{\partial H_\alpha}{\partial \boldsymbol{\pi}_i} \cdot \mathrm{d}\boldsymbol{\pi}_i \qquad \text{(vertical: heat)}. \tag{27}$$

The decomposition is exact on $\mathcal{E}$*:* $\mathrm{d}U$ *is exact by construction, but neither* δW *nor* δQ *is individually exact. The identity is purely kinematic—it is the chain rule applied to* $H_\alpha(\boldsymbol{\xi}, \boldsymbol{\pi})$ *and the horizontal–vertical decomposition of the tangent space* $T\mathcal{E} = \mathcal{H} \oplus \mathcal{V}$*; no equations of motion have been used.*

Proof. Equation (24) expresses $\mathrm{d}U$ as a sum of terms proportional to $\mathrm{d}\alpha_j$ (base directions) and terms proportional to $\mathrm{d}\boldsymbol{\xi}_i$, $\mathrm{d}\boldsymbol{\pi}_i$ (fiber directions). The product structure of the trivialized bundle provides a unique decomposition $T_e\mathcal{E} = \mathcal{H}_e \oplus \mathcal{V}_e$ at every point; collecting base components gives δW and fiber components gives δQ. □

The path-dependence of the work form

$$\delta W = \sum_j \mathcal{F}_j^{\text{tot}}\,\mathrm{d}\alpha_j$$

therefore arises entirely from the fiber-dependence of $\mathcal{F}_j^{\text{tot}}$: different paths in $\mathcal{B}$ sharing the same endpoints induce, via Newton's laws, different fiber trajectories and hence different values of the integrand. No geometric holonomy contributes.

The trivialization φ_α is not a modelling choice: it is determined by the physical apparatus that implements the boundary deformation. A rigid piston translating along one axis, an isotropic compression of all walls, and a flexible membrane changing shape non-uniformly all produce different maps $\Omega_0 \to \Omega(\alpha)$, even if they achieve the same macroscopic change $\Delta\alpha$. Correspondingly, the work–heat decomposition of Theorem V.7 depends on *how* the control parameters are varied, not merely on their initial and final values. This is not an ambiguity in the formalism but reflects a genuine physical fact: the split of energy transfer into work and heat depends on the operational definition of "boundary deformation," which is encoded in the trivialization.

Physical Content of the Decomposition

Definition V.8 (Total generalized force)**.** The *total generalized force conjugate to* α_j is

$$\mathcal{F}_j^{\rm tot} \;:=\; \frac{\partial H_\alpha}{\partial \alpha_j}\bigg|_{\boldsymbol{\xi},\,\boldsymbol{\pi}}, \tag{28}$$

so that the thermodynamic work takes the form $\delta W = \sum_j \mathcal{F}_j^{\rm tot}\,{\rm d}\alpha_j$.

Two consequences of the decomposition deserve to be stated explicitly, as they recover standard thermodynamic assertions from the geometry without independent postulation.

Proposition V.9 (Non-conservative forces produce only heat)**.** *Non-conservative forces do not contribute to* δW*: the total generalized force* $\mathcal{F}_j^{\rm tot}$ *depends only on the conservative sector of the dynamics.*

Proof. By Definition IV.2, the internal energy $U = K_{\rm rel} + \Phi$ absorbs all conservative work into the potential Φ; non-conservative forces are absent from U by construction. The fiber Hamiltonian $H_\alpha(\boldsymbol{\xi},\boldsymbol{\pi}) = K(\boldsymbol{\xi},\boldsymbol{\pi};\alpha) + \Phi_{\rm full}(\boldsymbol{\xi};\alpha)$ therefore contains no non-conservative contributions, and neither does its α-derivative $\mathcal{F}_j^{\rm tot} = \partial H_\alpha/\partial\alpha_j\big|_{\boldsymbol{\xi},\boldsymbol{\pi}}$. Non-conservative forces enter the equations of motion through the modified Hamilton equation $\dot{\boldsymbol{\pi}}_i = -\partial H_\alpha/\partial\boldsymbol{\xi}_i + \boldsymbol{f}_i^{\rm nc}$, which governs the fiber dynamics; they therefore affect only the vertical form δQ. ☐

This is the microscopic geometric origin of the thermodynamic dictum that dissipation produces heat, not work: it is not an independent postulate but a consequence of the fact that the work form is the horizontal projection of the differential of a Hamiltonian built from conservative forces alone.

Proposition V.10 (Heat vanishes for conservative systems at fixed boundary)**.** *Along a Hamiltonian trajectory at fixed control parameters* (${\rm d}\alpha = 0$)*, if all microscopic forces are conservative, then* $\delta Q = 0$.

Proof. At fixed α, Hamilton's equations on the fiber give $\dot H_\alpha = \{H_\alpha, H_\alpha\}_{\rm PB} = 0$. Since ${\rm d}\alpha = 0$ implies $\delta W = 0$, the first law ${\rm d}U = \delta W + \delta Q$ reduces to $\delta Q = {\rm d}H_\alpha = 0$. ☐

When non-conservative forces are present, the situation changes. Let $\boldsymbol{f}_i^{\rm nc} := J_\alpha(\boldsymbol{\xi}_i)^T \mathbf{F}_i^{\rm nc}$ be the canonical non-conservative force. The equations of motion become $\dot{\boldsymbol{\pi}}_i = -\partial H_\alpha/\partial\boldsymbol{\xi}_i + \boldsymbol{f}_i^{\rm nc}$ (with $\dot{\boldsymbol{\xi}}_i$ unchanged), and at fixed α:

$$\frac{\delta Q}{{\rm d}t}\bigg|_{\alpha\ \rm fixed} \;=\; \sum_i \frac{\partial H_\alpha}{\partial \boldsymbol{\pi}_i}\cdot \boldsymbol{f}_i^{\rm nc} \;=\; \sum_i \dot{\boldsymbol{\xi}}_i \cdot \boldsymbol{f}_i^{\rm nc}, \tag{29}$$

the power delivered by non-conservative forces through the internal degrees of freedom. Propositions V.9 and V.10 together give a complete microscopic characterisation: *heat is the energy transferred through the fiber by non-conservative forces; it vanishes identically when such forces are absent and the boundary is held fixed.*

Recovery of Standard Thermodynamic Expressions

As a consistency check, we verify that the formalism reproduces the standard equation of state for a gas by explicit computation of $\mathcal{F}_V^{\rm tot}$ in the simplest non-trivial case.

Consider N particles in a cubical box of volume $V = L^3$, with the affine scaling map $\varphi_V(\boldsymbol{\xi}) = V^{1/3}\boldsymbol{\xi}$, $\boldsymbol{\xi}_i \in [0,1]^3$. The Jacobian is $J_V = V^{1/3}\mathbf{I}$, so the metric (Definition V.4) is

$$g(V) \;=\; J_V^T J_V \;=\; V^{2/3}\,\mathbf{I}, \tag{30}$$

the canonical momenta (Definition V.3) are $\boldsymbol{\pi}_i = V^{1/3}\mathbf{p}_i'$, and the kinetic energy in canonical coordinates is

$$K \;=\; \sum_{i=1}^N \frac{\boldsymbol{\pi}_i^T\, g(V)^{-1}\,\boldsymbol{\pi}_i}{2m_i} \;=\; \sum_{i=1}^N \frac{|\boldsymbol{\pi}_i|^2}{2m_i\,V^{2/3}}. \tag{31}$$

a. Kinetic contribution. Differentiating (31) with respect to V at fixed $\boldsymbol{\pi}$:

$$\begin{aligned}\frac{\partial K}{\partial V}\bigg|_{\boldsymbol{\pi}} &= \sum_i \frac{|\boldsymbol{\pi}_i|^2}{2m_i}\cdot\frac{\partial}{\partial V}\Big(V^{-2/3}\Big)\\ &= -\frac{2}{3}\,V^{-5/3}\sum_i \frac{|\boldsymbol{\pi}_i|^2}{2m_i}\\ &= -\frac{2}{3V}\,K.\end{aligned} \tag{32}$$

b. Potential contribution. For pairwise interactions $\Phi(\mathbf{r}') = \sum_{i<k}\phi(|\mathbf{r}_i' - \mathbf{r}_k'|)$, the potential in canonical coordinates is

$$\Phi_{\rm full}(\boldsymbol{\xi};V) \;=\; \sum_{i<k}\phi\big(V^{1/3}|\boldsymbol{\xi}_i - \boldsymbol{\xi}_k|\big) \;=\; \sum_{i<k}\phi(r_{ik}), \tag{33}$$

where $r_{ik} := |\mathbf{r}'_i - \mathbf{r}'_k| = V^{1/3}|\boldsymbol{\xi}_i - \boldsymbol{\xi}_k|$. At fixed $\boldsymbol{\xi}$:

$$\left.\frac{\partial r_{ik}}{\partial V}\right|_{\boldsymbol{\xi}} = \frac{1}{3}V^{-2/3}|\boldsymbol{\xi}_i - \boldsymbol{\xi}_k| = \frac{r_{ik}}{3V}. \quad (34)$$

Applying the chain rule to (33):

$$\left.\frac{\partial \Phi_{\text{full}}}{\partial V}\right|_{\boldsymbol{\xi}} = \sum_{i<k}\phi'(r_{ik})\left.\frac{\partial r_{ik}}{\partial V}\right|_{\boldsymbol{\xi}} = \frac{1}{3V}\sum_{i<k} r_{ik}\,\phi'(r_{ik}). \quad (35)$$

c. Total generalized force. Combining (32) and (35) via Definition V.8:

$$\mathcal{F}_V^{\text{tot}} = \left.\frac{\partial H_V}{\partial V}\right|_{\boldsymbol{\xi},\boldsymbol{\pi}} = \underbrace{-\frac{2}{3V}K}_{\text{kinetic}} + \underbrace{\frac{1}{3V}\sum_{i<k} r_{ik}\,\phi'(r_{ik})}_{\text{interaction virial}}. \quad (36)$$

Equation (36) is the microscopic generalized force for a single microstate; it is the integrand of the virial theorem. The passage to the macroscopic pressure $P = -\bar{\mathcal{F}}_V^{\text{tot}}$ requires fiber averaging over an equilibrium measure, which is the subject of Section VI.

VI. FROM MICROSCOPIC TO MACROSCOPIC: FIBER AVERAGING AND THE MACROSCOPIC FIRST LAW

The work–heat decomposition of Theorem V.7,

$$\mathrm{d}U = \delta W + \delta Q, \qquad \delta W = \sum_j \mathcal{F}_j^{\text{tot}}\,\mathrm{d}\alpha_j, \quad (37)$$

is an exact identity on the total phase space $\mathcal{E}$, holding at every point $(\alpha, \boldsymbol{\xi}, \boldsymbol{\pi})$. Every quantity appearing in it—the internal energy $U = H_\alpha(\boldsymbol{\xi}, \boldsymbol{\pi})$, the generalized force $\mathcal{F}_j^{\text{tot}} = \partial H_\alpha/\partial\alpha_j|_{\boldsymbol{\xi},\boldsymbol{\pi}}$, and the heat form δQ—depends on the full microstate $(\boldsymbol{\xi}, \boldsymbol{\pi}) \in \mathcal{F}_\alpha$.

A macroscopic observer has access to the control parameters $\alpha \in \mathcal{B}$ and to coarse thermodynamic variables such as pressure and temperature; the instantaneous microstate is inaccessible. Macroscopic thermodynamics must therefore be formulated in terms of functions on $\mathcal{B}$ alone. This creates a precise mathematical problem: *how does one systematically eliminate the fiber-dependence of the quantities in (37) to obtain an identity involving only functions of α?*

The Projection Problem

A function $f : \mathcal{E} \to \mathbb{R}$ that is not constant on fibers does not descend to a function on $\mathcal{B}$. The passage from one to the other requires a procedure that assigns, to each function f on $\mathcal{E}$ and each fiber $\mathcal{F}_\alpha$, a single real number—the "macroscopic value" of f at α. We now show that the requirements of linearity, positivity, and normalisation uniquely force this procedure to be integration against a probability measure.

Proposition VI.1 (Geometric necessity of the fiber measure)**.** *For each $\alpha \in \mathcal{B}$, let $\Lambda_\alpha : C_c^\infty(\mathcal{F}_\alpha) \to \mathbb{R}$ be a linear functional satisfying:*

(i) positivity*: $f \geq 0 \Rightarrow \Lambda_\alpha(f) \geq 0$;*
(ii) normalisation*: $\sup\{\Lambda_\alpha(f) : f \in C_c^\infty(\mathcal{F}_\alpha),\ 0 \leq f \leq 1\} = 1$.*

Then there exists a unique Radon probability measure μ_α on $\mathcal{F}_\alpha$ such that

$$\Lambda_\alpha(f) = \int_{\mathcal{F}_\alpha} f\,\mathrm{d}\mu_\alpha \qquad \text{for all } f \in C_c^\infty(\mathcal{F}_\alpha). \quad (38)$$

Remark VI.2. The requirements on Λ_α are dictated by elementary physical consistency. *Linearity*: if two microscopic observables f and g have macroscopic values $\Lambda_\alpha(f)$ and $\Lambda_\alpha(g)$, then their sum $f + g$—measured on the same microstate—must have macroscopic value $\Lambda_\alpha(f) + \Lambda_\alpha(g)$; energy, for instance, is additive. *Positivity*: a microscopic observable that is everywhere non-negative (such as the kinetic energy) must have a non-negative macroscopic value. *Normalisation*: the macroscopic value of a constant observable equals that constant; in particular, $\Lambda_\alpha(1) = 1$.

Proof. The phase-space fiber $\mathcal{F}_\alpha = T^*\mathcal{Q}_\alpha$ is a locally compact Hausdorff space. By condition (i), Λ_α is a positive linear functional on $C_c(\mathcal{F}_\alpha)$ (since C_c^∞ is dense in C_c and positivity extends by continuity). The Riesz–Markov–Kakutani representation theorem [47] yields a unique Radon measure μ_α satisfying (38). Condition (ii) ensures $\mu_\alpha(\mathcal{F}_\alpha) = 1$: for any exhaustion $K_1 \subset K_2 \subset \cdots$ of $\mathcal{F}_\alpha$ by compact sets, choose $f_n \in C_c^\infty$ with $0 \leq f_n \leq 1$ and $f_n = 1$ on K_n; then $\mu_\alpha(\mathcal{F}_\alpha) = \lim_n \int f_n\,\mathrm{d}\mu_\alpha = \lim_n \Lambda_\alpha(f_n) \leq 1$, and the supremum condition forces equality. □

Proposition VI.1 establishes that the projection takes the form of integration against a Radon probability measure μ_α on $\mathcal{F}_\alpha$. For compactly supported observables, every such measure suffices. But the physically relevant observables—the generalized force $\mathcal{F}_j^{\text{tot}}$, the kinetic energy K, the fiber Hamiltonian H_α—are polynomial in the momenta $\boldsymbol{\pi}_i$ and therefore unbounded on the non-compact fiber $\mathcal{F}_\alpha$. For the fiber average $\bar{f}(\alpha) := \int_{\mathcal{F}_\alpha} f\,\mathrm{d}\mu_\alpha$ to be well-defined, the measure must satisfy

$$\int_{\mathcal{F}_\alpha} |\boldsymbol{\pi}|^{2k}\,\mathrm{d}\mu_\alpha < \infty \qquad \text{for all } k \in \mathbb{N}, \quad (39)$$

i.e., all polynomial moments in the momenta must be finite. This is a non-trivial constraint: it excludes, for instance, any measure that is uniform in the momenta, and requires μ_α to decay sufficiently rapidly as $|\boldsymbol{\pi}| \to \infty$.

The bundle geometry therefore determines three things: that a probability measure is needed, that it must live on the fiber, and that it must have finite moments of all orders. What it does not determine is the specific functional form of μ_α. That selection requires physical input, and it is here—and only here—that the standard apparatus of equilibrium statistical mechanics enters the framework.

Selecting the Measure: From Symplectic Geometry to the Microcanonical Ensemble

The projection problem requires a Radon probability measure μ_α on each fiber $\mathcal{F}_\alpha$ with finite moments of all orders (39). We now show that three physically motivated constraints—each stated precisely—narrow the class of admissible measures, and that a single additional hypothesis (ergodicity) selects one uniquely.

a. The canonical volume form. Each fiber $\mathcal{F}_\alpha = T^*\mathcal{Q}_\alpha$ is a symplectic manifold with the canonical symplectic form $\omega_{\rm symp} = \sum_i \mathrm{d}\boldsymbol{\xi}_i \wedge \mathrm{d}\boldsymbol{\pi}_i$ (Definition V.3). The top exterior power of $\omega_{\rm symp}$ defines a volume form—the *Liouville measure*—on each fiber:

$$\mathrm{d}\Gamma_\alpha \;\coloneqq\; \frac{\omega_{\rm symp}^{\wedge(3N-3)}}{(3N-3)!} \;=\; {\prod_i}' \mathrm{d}^3\boldsymbol{\xi}_i\, \mathrm{d}^3\boldsymbol{\pi}_i, \tag{40}$$

where the primed product denotes the volume form on the $(6N-6)$-dimensional constraint surface $\sum_i m_i \mathbf{r}'_i = \mathbf{0}$, $\sum_i m_i \mathbf{v}'_i = \mathbf{0}$.

Two properties of $\mathrm{d}\Gamma_\alpha$ are immediate:

(a) *Intrinsic character.* $\mathrm{d}\Gamma_\alpha$ is determined by the symplectic structure alone; it requires no metric, no Riemannian structure, and no physical input beyond the phase-space geometry.

(b) *α-independence in canonical coordinates.* In the Darboux coordinates $(\boldsymbol{\xi}, \boldsymbol{\pi})$, the expression (40) is manifestly independent of α. The α-dependence that would appear in the physical coordinates $(\mathbf{r}', \mathbf{p}')$—through the Jacobian $\det(J_\alpha)^N$ of the trivialization—is absorbed into the canonical momenta $\boldsymbol{\pi}_i = J_\alpha^T \mathbf{p}'_i$. This simplifies all subsequent phase-space integrals.

The Liouville measure is an infinite measure on $\mathcal{F}_\alpha$ (the momentum directions are unbounded) and therefore cannot serve as a probability measure. It provides the *reference measure*—the analogue of Lebesgue measure on $\mathbb{R}^n$—against which the probability measure μ_α must be absolutely continuous.

b. Restriction to the energy shell. For an isolated system with no non-conservative forces, the fiber Hamiltonian H_α is conserved at fixed α (Proposition V.10). The trajectory at energy U is therefore confined to the energy shell

$$\Sigma_\alpha(U) \;\coloneqq\; \big\{(\boldsymbol{\xi}, \boldsymbol{\pi}) \in \mathcal{F}_\alpha : H_\alpha(\boldsymbol{\xi}, \boldsymbol{\pi}) = U\big\}. \tag{41}$$

Any probability measure intended to describe the equilibrium macrostate at energy U must be supported on $\Sigma_\alpha(U)$: a measure that assigns positive weight to states with $H_\alpha \neq U$ would describe a system at the wrong energy.

The energy shell $\Sigma_\alpha(U)$ is a $(6N-7)$-dimensional submanifold of $\mathcal{F}_\alpha$ whenever U is a regular value of H_α (i.e., $\mathrm{d}H_\alpha \neq 0$ on $\Sigma_\alpha(U)$). When the potential Φ is confining—$\Phi(\mathbf{r}') \to +\infty$ as any particle approaches the boundary of $\Omega(\alpha)$—the kinetic energy is bounded above by $U - \Phi$, and $\Sigma_\alpha(U)$ is compact. Compactness guarantees that all polynomial moments are finite, so the moment condition (39) is automatically satisfied for any probability measure supported on $\Sigma_\alpha(U)$.

For hard-wall confinement (infinite potential at the boundary, zero in the interior), the potential energy is bounded ($0 \leq \Phi \leq U$ for pairwise potentials with a hard core), and the energy shell is again compact. For systems with purely kinetic energy ($\Phi = 0$ in the interior, as in an ideal gas), $\Sigma_\alpha(U) = \{\sum_i |\boldsymbol{\pi}_i|^2/(2m_i g_\alpha^{-1}) = U\}$ is an ellipsoid in momentum space—compact.

c. Flow-invariance. Macroscopic equilibrium means, by definition, that macroscopic observables do not change with time. If the fiber average $\bar{f}(\alpha) = \int f\, \mathrm{d}\mu_\alpha$ is to be time-independent for every observable f, the measure μ_α must be invariant under the Hamiltonian flow $\phi^t_{H_\alpha}$ on the fiber:

$$(\phi^t_{H_\alpha})_* \mu_\alpha \;=\; \mu_\alpha \qquad \text{for all } t. \tag{42}$$

By Liouville's theorem [4], the Hamiltonian flow preserves the Liouville measure $\mathrm{d}\Gamma_\alpha$. However, this does *not* imply that every measure absolutely continuous with respect to $\mathrm{d}\Gamma_\alpha$ is flow-invariant: a measure $\rho\, \mathrm{d}\Gamma_\alpha$ is flow-invariant if and only if ρ is a constant of the motion ($\rho \circ \phi^t_{H_\alpha} = \rho$ for all t). Flow-invariance is therefore an independent constraint, not a consequence of absolute continuity.

d. Absolute continuity. We require $\mu_\alpha \ll \mathrm{d}\Gamma_\alpha|_{\Sigma_\alpha(U)}$: the probability measure must be absolutely continuous with respect to the Liouville measure restricted to the energy shell. This excludes singular measures concentrated on individual orbits or invariant tori, which would describe a system prepared in a specific microscopic trajectory rather than a macroscopic equilibrium state. Physically, this is the requirement that the macroscopic description does not smuggle in microscopic information through the choice of measure.

The four constraints are:

(C1) μ_α is a probability measure on $\mathcal{F}_\alpha$;
(C2) μ_α is supported on the energy shell $\Sigma_\alpha(U)$;
(C3) μ_α is invariant under the Hamiltonian flow of H_α;
(C4) μ_α is absolutely continuous w.r.t. $\mathrm{d}\Gamma_\alpha|_{\Sigma_\alpha(U)}$.

These constraints do not yet determine μ_α uniquely: if the Hamiltonian flow on $\Sigma_\alpha(U)$ admits non-trivial invariant subsets of positive Liouville measure, then any absolutely continuous measure supported on such a subset satisfies (C1)–(C4). Uniqueness requires one further hypothesis.

Assumption VI.3 (Ergodicity)**.** The Hamiltonian flow of H_α on the energy shell $\Sigma_\alpha(U)$ is *ergodic* with respect to the Liouville measure: every flow-invariant measurable subset of $\Sigma_\alpha(U)$ has either zero or full Liouville measure.

Ergodicity is a hypothesis about the dynamics, not a consequence of the symplectic geometry. It holds rigorously for billiard systems in convex domains [51] and for certain classes of geodesic flows on negatively curved manifolds [3], and is believed to hold for generic many-body systems with short-range interactions at sufficiently high energy [38]. It fails for integrable systems (e.g., a collection of non-interacting harmonic oscillators), where the phase space is foliated by invariant tori and the microcanonical measure is not the unique flow-invariant probability measure on the energy shell.

Proposition VI.4 (Uniqueness of the microcanonical measure)**.** *Under Assumption VI.3, the unique probability measure on* $\Sigma_\alpha(U)$ *satisfying* (C1)–(C4) *is*

$$\mathrm{d}\mu_\alpha^{\mathrm{mc}} \;=\; \frac{1}{\Sigma(\alpha,U)}\,\delta\big(H_\alpha-U\big)\,\mathrm{d}\Gamma_\alpha, \tag{43}$$

where

$$\Sigma(\alpha,U) \;\coloneqq\; \int_{\mathcal{F}_\alpha}\delta\big(H_\alpha-U\big)\,\mathrm{d}\Gamma_\alpha \tag{44}$$

is the density of states*—the Liouville surface area of the energy shell.*

Proof. By (C2) and (C4), μ_α has a density ρ with respect to $\mathrm{d}\Gamma_\alpha|_{\Sigma_\alpha(U)}$: $\mathrm{d}\mu_\alpha=\rho\,\mathrm{d}\Gamma_\alpha|_{\Sigma_\alpha(U)}$. By (C3), ρ must be invariant under the Hamiltonian flow. By Assumption VI.3, the only flow-invariant L^1 functions on $\Sigma_\alpha(U)$ are constant almost everywhere with respect to the Liouville measure. Therefore ρ is constant $\mathrm{d}\Gamma_\alpha$-almost everywhere on $\Sigma_\alpha(U)$. By the co-area formula, $\int_{\Sigma_\alpha(U)}\mathrm{d}\Gamma_\alpha|_{\Sigma_\alpha(U)}=\Sigma(\alpha,U)$, so the normalization condition (C1) fixes $\rho=1/\Sigma(\alpha,U)$. □

The microcanonical measure is therefore not postulated; it is the unique measure compatible with the four constraints (C1)–(C4) and the ergodic hypothesis. The traditional "principle of equal a priori probability"—the assertion that all microstates on the energy shell are equally likely—is here a *theorem*: it follows from the flow-invariance of the density ρ and the ergodicity of the dynamics, not from an independent probabilistic axiom.

Definition VI.5 (Fiber average)**.** For any observable $f:\mathcal{E}\to\mathbb{R}$, the *fiber average* at (α,U) is

$$\begin{aligned}\bar f(\alpha,U) &\coloneqq \int_{\mathcal{F}_\alpha} f\,\mathrm{d}\mu_\alpha^{\mathrm{mc}}\\ &= \frac{1}{\Sigma(\alpha,U)}\int_{\mathcal{F}_\alpha} f\,\delta(H_\alpha-U)\,\mathrm{d}\Gamma_\alpha.\end{aligned} \tag{45}$$

This defines a smooth function $\bar f:\mathcal{B}\times\mathbb{R}\to\mathbb{R}$ whenever U is a regular value of H_α and $\Sigma_\alpha(U)$ is compact.

Remark VI.6 (Alternatives to the ergodic hypothesis)*.* If Assumption VI.3 is dropped, the microcanonical measure can be selected by alternative principles. The maximum-entropy approach [30] selects $\rho=1/\Sigma$ as the unique maximiser of the Gibbs entropy $S_G[\mu]=-k_B\int\rho\ln\rho\,\mathrm{d}\Gamma_\alpha$ over all absolutely continuous probability measures on $\Sigma_\alpha(U)$; this replaces a dynamical hypothesis with an information-theoretic one. In the thermodynamic limit $N\to\infty$, the equivalence of ensembles [48, 52] ensures that macroscopic predictions are insensitive to the choice of measure within the class (C1)–(C4), rendering the uniqueness question moot for bulk observables. The present framework is compatible with any of these justifications; Assumption VI.3 is adopted because it is the only one that is purely dynamical and valid at finite N.

Derived Ensembles: The Canonical and Grand Canonical Measures

The microcanonical measure describes an isolated system at fixed energy. In practice, thermodynamic systems are rarely isolated: they exchange energy with a heat bath, or energy and particles with a reservoir. We now show that each standard ensemble arises from the microcanonical by a single mathematical operation—marginalisation over unobserved degrees of freedom—applied to progressively larger composite systems. No new postulates are introduced; each measure is a *theorem* of the microcanonical framework.

e. The composite system. Consider a system S in thermal contact with a heat bath B. Both are described by the framework of the preceding sections: S has fiber $\mathcal{F}_\alpha^S$ with Hamiltonian H_α^S, and B has fiber $\mathcal{F}^B$ with Hamiltonian H^B. The composite system $S+B$ has the product fiber $\mathcal{F}_\alpha^{S+B}=\mathcal{F}_\alpha^S\times\mathcal{F}^B$ and the total Hamiltonian

$$\begin{aligned}H_\alpha^{S+B}(\boldsymbol{\xi}^S,\boldsymbol{\pi}^S,\boldsymbol{\xi}^B,\boldsymbol{\pi}^B) &= H_\alpha^S(\boldsymbol{\xi}^S,\boldsymbol{\pi}^S)\\ &\quad + H^B(\boldsymbol{\xi}^B,\boldsymbol{\pi}^B)+H^{\mathrm{int}},\end{aligned} \tag{46}$$

where H^{int} is the interaction between S and B.

f. The weak-coupling assumption. We assume that the interaction energy is negligible compared to the energies of S and B separately:

$$|H^{\mathrm{int}}| \;\ll\; H_\alpha^S,\quad |H^{\mathrm{int}}| \;\ll\; H^B. \tag{47}$$

Under this assumption, the total energy constraint becomes $H_\alpha^S+H^B=U_{\mathrm{tot}}$, and the Liouville measure on $\mathcal{F}_\alpha^{S+B}$ factorises: $\mathrm{d}\Gamma^{S+B}=\mathrm{d}\Gamma^S\cdot\mathrm{d}\Gamma^B$.

g. Marginalisation. The composite system is isolated at total energy U_{tot}, so it is described by the microcanonical measure on $\Sigma_\alpha^{S+B}(U_{\mathrm{tot}})$. The macroscopic observer of S has no access to the bath degrees of freedom $(\boldsymbol{\xi}^B,\boldsymbol{\pi}^B)$. The appropriate measure on $\mathcal{F}_\alpha^S$ is therefore the *marginal* of the microcanonical measure, obtained by integrating out the bath:

$$\mathrm{d}\mu_\alpha^S \;\propto\; \left[\int_{\mathcal{F}^B}\delta\big(H_\alpha^S+H^B-U_{\mathrm{tot}}\big)\,\mathrm{d}\Gamma^B\right]\mathrm{d}\Gamma^S. \tag{48}$$

The quantity in brackets is the density of states of the bath at energy $U_{\mathrm{tot}}-H_\alpha^S$:

$$\Omega_B(U_{\mathrm{tot}}-H_\alpha^S) \;\coloneqq\; \int_{\mathcal{F}^B}\delta\big(H^B-(U_{\mathrm{tot}}-H_\alpha^S)\big)\,\mathrm{d}\Gamma^B. \tag{49}$$

h. The thermodynamic limit of the bath. Define the bath entropy $S_B(E) := k_B \ln \Omega_B(E)$. Since the bath is macroscopic ($N_B \gg N_S$) and $H^S_\alpha \ll U_{\rm tot}$, we expand S_B to first order in H^S_α:

$$\begin{aligned} & S_B(U_{\rm tot} - H^S_\alpha) \\ =\ & S_B(U_{\rm tot}) \ - \ H^S_\alpha \frac{\partial S_B}{\partial E}\bigg|_{U_{\rm tot}} \ + \ \mathcal{O}\bigg(\frac{H^S_\alpha}{U_{\rm tot}}\bigg)^2 . \end{aligned} \tag{50}$$

The derivative $\partial S_B/\partial E =: 1/T$ defines the temperature of the bath. Exponentiating:

$$\begin{aligned} & \Omega_B(U_{\rm tot} - H^S_\alpha) \\ =\ & \Omega_B(U_{\rm tot}) \exp\bigg(-\frac{H^S_\alpha}{k_B T}\bigg) \bigg[1 + \mathcal{O}\bigg(\frac{H^S_\alpha}{U_{\rm tot}}\bigg)^2\bigg] . \end{aligned} \tag{51}$$

The error term is of order N_S/N_B and vanishes in the limit $N_B \to \infty$ at fixed N_S.

Proposition VI.7 (The canonical measure from marginalisation)**.** *Under the weak-coupling assumption (47) and in the thermodynamic limit of the bath ($N_B \to \infty$ at fixed T), the marginal measure on $\mathcal{F}^S_\alpha$ is*

$$\begin{aligned} d\mu^{\rm can}_\alpha &= \frac{1}{Z(\alpha,\beta)}\, e^{-\beta H^S_\alpha(\boldsymbol{\xi}^S,\boldsymbol{\pi}^S)}\, d\Gamma^S , \\ Z(\alpha,\beta) &:= \int_{\mathcal{F}^S_\alpha} e^{-\beta H^S_\alpha}\, d\Gamma^S , \end{aligned} \tag{52}$$

where $\beta := 1/(k_B T)$ and $Z(\alpha,\beta)$ is the canonical partition function*.*

Proof. Substituting (51) into (48) and dropping the $\mathcal{O}(N_S/N_B)$ correction:

$$d\mu^S_\alpha \ \propto\ e^{-\beta H^S_\alpha}\, d\Gamma^S . \tag{53}$$

Normalisation to unit total mass gives

$$Z(\alpha,\beta) = \int e^{-\beta H^S_\alpha}\, d\Gamma^S \ .$$

The Gaussian decay $e^{-\beta|\boldsymbol{\pi}|^2/(2m)}$ in the momenta guarantees that $Z < \infty$ and that all polynomial moments are finite, verifying the moment condition (39). □

The canonical measure is therefore not postulated; it is the marginal of the microcanonical measure on the composite system, in the thermodynamic limit of the bath. The Boltzmann factor $e^{-\beta H}$ is a consequence of the logarithmic dependence of the bath entropy on energy—which is itself a property of the density of states $\Omega_B(E)$—not an independent axiom.

Remark VI.8 (Moment condition for the canonical measure)*.* The exponential decay $e^{-\beta|\boldsymbol{\pi}|^2/(2m)}$ in the momenta guarantees $\int |\boldsymbol{\pi}|^{2k}\, d\mu^{\rm can}_\alpha < \infty$ for all $k \in \mathbb{N}$, so the canonical measure satisfies the moment condition (39) on the full (non-compact) fiber $\mathcal{F}^S_\alpha$ without requiring restriction to an energy shell.

i. The grand canonical measure: exchanging both energy and particles. The canonical measure describes a system that exchanges energy with a reservoir but has a fixed number of particles. When the system also exchanges particles with the reservoir, the particle number N_S becomes a fluctuating quantity, and the appropriate fiber changes.

j. The extended fiber. For each particle number N, the system has a fiber $\mathcal{F}^{S,N}_\alpha$ with Hamiltonian $H^{S,N}_\alpha$. The *extended fiber* is the disjoint union over all admissible particle numbers:

$$\widehat{\mathcal{F}}^S_\alpha \ :=\ \bigsqcup_{N=0}^{\infty} \mathcal{F}^{S,N}_\alpha . \tag{54}$$

An observable on $\widehat{\mathcal{F}}^S_\alpha$ is a family of functions $\{f_N\}_{N\ge 0}$ with $f_N : \mathcal{F}^{S,N}_\alpha \to \mathbb{R}$.

k. Marginalisation over the reservoir. The composite system $S+B$ is isolated at fixed total energy $U_{\rm tot}$ and total particle number $N_{\rm tot} = N_S + N_B$. The microcanonical measure is defined on the composite energy shell with the constraint $N_S + N_B = N_{\rm tot}$. Marginalising over both the bath phase-space coordinates *and* the bath particle number $N_B = N_{\rm tot} - N_S$:

$$d\mu^S_\alpha\big|_{N_S} \ \propto\ \Omega_B(U_{\rm tot} - H^{S,N_S}_\alpha,\ N_{\rm tot} - N_S)\, d\Gamma^{S,N_S}, \tag{55}$$

where $\Omega_B(E,N)$ is the bath density of states at energy E and particle number N.

l. Thermodynamic limit of the bath. Expanding the bath entropy $S_B(E,N) = k_B \ln \Omega_B(E,N)$ to first order in both H^{S,N_S}_α and N_S:

$$\begin{aligned} S_B(U_{\rm tot} - H^{S,N_S}_\alpha, N_{\rm tot} - N_S) = S_B(U_{\rm tot}, N_{\rm tot}) - \frac{H^{S,N_S}_\alpha}{T} \\ + \frac{\mu\, N_S}{T} + \mathcal{O}\bigg(\frac{1}{N_B}\bigg), \end{aligned} \tag{56}$$

where

$$\frac{1}{T} \ :=\ \frac{\partial S_B}{\partial E}\bigg|_N , \qquad \frac{\mu}{T} \ :=\ -\frac{\partial S_B}{\partial N}\bigg|_E \tag{57}$$

define the temperature and chemical potential of the bath.

Proposition VI.9 (The grand canonical measure from marginalisation)**.** *Under the weak-coupling assumption and in the thermodynamic limit of the bath ($N_B \to \infty$ at fixed T and μ), the marginal measure on the extended fiber $\widehat{\mathcal{F}}^S_\alpha$ is*

$$d\mu^{\rm gc}_\alpha = \frac{1}{\mathcal{Z}(\alpha,\beta,\mu)} \sum_{N=0}^{\infty} \frac{e^{\beta\mu N}}{N!\, h^{3N}}\, e^{-\beta H^{S,N}_\alpha}\, d\Gamma^{S,N} \tag{58}$$

where

$$\mathcal{Z}(\alpha,\beta,\mu) := \sum_{N=0}^{\infty} (N!\, h^{3N})^{-1}\, e^{\beta\mu N} \int e^{-\beta H^{S,N}_\alpha}\, d\Gamma^{S,N}$$

is the grand canonical partition function*, and the factors* $N!$ *and* h^{3N} *account for the indistinguishability of identical particles and the semiclassical phase-space cell, respectively.*

Proof. Exponentiating (56) and substituting into (55), the marginal measure at particle number N_S is

$$\mathrm{d}\mu_\alpha^S\big|_{N_S} \;\propto\; e^{\beta\mu N_S}\, e^{-\beta H_\alpha^{S,N_S}}\, \mathrm{d}\Gamma^{S,N_S}. \qquad (59)$$

Summing over N_S and normalising gives the classical grand canonical measure (58) without the combinatorial factors; the factors $N!$ and h^{3N} are inserted by hand to match the standard notation, as discussed in Remark VI.10. The convergence of $\mathcal{Z}$ for μ below the condensation threshold is guaranteed by the exponential decay in $H_\alpha^{S,N}$ at each N and the factorial suppression at large N. □

Remark VI.10 (Quantum origin of the combinatorial factors)*.* The factors $N!$ and h^{3N} in (58) cannot be derived within the classical framework of this paper. The factor h^{3N} sets the phase-space cell size and originates from the semiclassical limit of quantum mechanics; the factor $N!$ corrects for the quantum indistinguishability of identical particles, resolving the Gibbs paradox [21]. In the classical theory, these factors affect only the absolute value of the entropy (through $\ln Z$) and not the equations of state or the thermodynamic forces; they are included here to make contact with the standard notation. A fully classical treatment would omit them, at the cost of an undetermined additive constant in S.

Remark VI.11 (The logical hierarchy of ensembles)*.* Each step is the same mathematical operation—integration of the microcanonical measure over unobserved degrees of freedom—applied to a composite system in which the subsystem exchanges progressively more quantities with the reservoir. No new postulate enters at any stage: the Boltzmann factor $e^{-\beta H}$ and the fugacity $e^{\beta\mu N}$ are consequences of the logarithmic structure of the bath entropy, and the microcanonical measure itself is the unique measure compatible with (C1)–(C4) and ergodicity. The entire apparatus of equilibrium statistical mechanics is, within this framework, a consequence of Newtonian mechanics, the fiber-bundle geometry, and the ergodic hypothesis.

The Macroscopic Work Form

Applying the fiber average (Definition VI.5) to the microscopic generalized force (28) yields a function on $\mathcal{B}$:

Definition VI.12 (Thermodynamic generalized force)**.**

$$\bar{\mathcal{F}}_j^{\,\mathrm{tot}}(\alpha,U) \;:=\; \int_{\mathcal{F}_\alpha} \mathcal{F}_j^{\mathrm{tot}}\,\mathrm{d}\mu_\alpha^{\mathrm{mc}} \;=\; \left\langle \frac{\partial H_\alpha}{\partial \alpha_j}\bigg|_{\boldsymbol{\xi},\boldsymbol{\pi}} \right\rangle_{\mu_\alpha^{\mathrm{mc}}}. \qquad (60)$$

The *macroscopic work form* on $\mathcal{B}$ is

$$\delta\bar{W} \;:=\; \sum_j \bar{\mathcal{F}}_j^{\,\mathrm{tot}}\,\mathrm{d}\alpha_j. \qquad (61)$$

The microscopic work form $\delta W = \sum_j \mathcal{F}_j^{\mathrm{tot}}\,\mathrm{d}\alpha_j$ (Theorem V.7) is a 1-form on the total space $\mathcal{E}$; its coefficient depends on the full microstate and fluctuates along every trajectory. The macroscopic form $\delta\bar{W}$ is a 1-form on the base $\mathcal{B}$; its coefficient $\bar{\mathcal{F}}_j^{\,\mathrm{tot}}$ depends on the thermodynamic state (α,U) alone. These are different mathematical objects on different spaces.

Recovery of the Virial Equation of State

With the fiber average now defined, we complete the computation deferred from §V. Recall that for N particles in a cubical box of volume $V = L^3$ with the affine scaling map $\varphi_V(\boldsymbol{\xi}) = V^{1/3}\boldsymbol{\xi}$, the microscopic generalized force conjugate to V was found to be (equation (36)):

$$\mathcal{F}_V^{\mathrm{tot}} \;=\; -\frac{2}{3V}\,K \;+\; \frac{1}{3V}\sum_{i<k} r_{ik}\,\phi'(r_{ik}). \qquad (62)$$

This is a function on the total space $\mathcal{E}$: it depends on the microstate $(\boldsymbol{\xi},\boldsymbol{\pi})$ through K and through the interparticle separations r_{ik}. We now apply Definition VI.5 to obtain a function on $\mathcal{B}$.

m. The thermodynamic pressure. The macroscopic work form (61) for the single control parameter $\alpha = V$ is $\delta\bar{W} = \bar{\mathcal{F}}_V^{\,\mathrm{tot}}\,\mathrm{d}V$. The thermodynamic pressure is defined as the negative of the averaged generalized force:

$$P \;:=\; -\bar{\mathcal{F}}_V^{\,\mathrm{tot}} \;=\; -\int_{\mathcal{F}_V} \mathcal{F}_V^{\mathrm{tot}}\,\mathrm{d}\mu_V, \qquad (63)$$

so that $\delta\bar{W} = -P\,\mathrm{d}V$, recovering the standard thermodynamic sign convention.

n. Fiber averaging of the kinetic term. The kinetic energy $K = \sum_i |\boldsymbol{\pi}_i|^2/(2m_i V^{2/3})$ is quadratic in the momenta. For any system with kinetic energy quadratic in the canonical momenta, the microcanonical equipartition theorem [29, 31] gives

$$\left\langle \pi_{i,a}\,\frac{\partial H_\alpha}{\partial \pi_{i,a}} \right\rangle_{\mathrm{mc}} \;=\; k_B T \qquad (64)$$

for each momentum component $\pi_{i,a}$, where T is defined by $1/T = \partial S/\partial U|_\alpha = k_B\Sigma/\Omega$. Since $\partial H_\alpha/\partial \pi_{i,a} = \pi_{i,a}/(m_i V^{2/3})$ for the cubical-box Hamiltonian, summing over all $3(N-1)$ independent momentum components yields

$$\langle K\rangle \;=\; \frac{3(N-1)}{2}\,k_B T \;\approx\; \frac{3}{2}\,N k_B T, \qquad (65)$$

where the approximate equality holds for $N \gg 1$. This result is valid for interacting systems: the equipartition theorem applies to each quadratic momentum degree of freedom independently of the potential energy.

o. Fiber averaging of the interaction term. The virial sum $\sum_{i<k} r_{ik}\,\phi'(r_{ik})$ depends on positions alone. Its fiber average $\langle\sum_{i<k} r_{ik}\,\phi'(r_{ik})\rangle$ is a function of (V,U)—or equivalently (V,T)—that encodes all information about inter-particle correlations at equilibrium.

p. The macroscopic equation of state. Substituting into (63):

$$P \;=\; \frac{2\langle K\rangle}{3V} \;-\; \frac{1}{3V}\left\langle \sum_{i<k} r_{ik}\,\phi'(r_{ik})\right\rangle. \tag{66}$$

Using (65):

$$P \;=\; \frac{Nk_BT}{V} \;-\; \frac{1}{3V}\left\langle \sum_{i<k} r_{ik}\,\phi'(r_{ik})\right\rangle. \tag{67}$$

The first term is the ideal gas law; the second is the interaction correction. Together they reproduce the virial equation of state [13].

q. Consistency checks. Three limiting cases confirm that the formalism is internally consistent:

(i) *Ideal gas* ($\phi = 0$). The interaction term vanishes, and $P = Nk_BT/V$: the entire pressure is of kinetic origin, arising from the α-dependence of the fiber metric $g(V) = V^{2/3}\mathbf{I}$ in the kinetic energy. This confirms that the formalism correctly accounts for the work done by a confined non-interacting gas—a non-trivial check, since the pressure arises not from particle–wall collisions (which are absent in the fiber-averaged description) but from the geometric fact that the fiber metric depends on V.

(ii) *Hard spheres* ($\phi = 0$ for $r > \sigma$, $\phi = +\infty$ for $r < \sigma$). The virial sum reduces to a contact-value contribution, and (67) reproduces the hard-sphere virial series whose leading term is $P = (Nk_BT/V)(1 + Nb/(V))$ with $b = \frac{2}{3}\pi\sigma^3$ [24].

(iii) *Macroscopic first law.* Substituting $\delta\bar{W} = -P\,\mathrm{d}V$ into the macroscopic first law (69): $\mathrm{d}U = -P\,\mathrm{d}V + \delta\bar{Q}$, which is the standard thermodynamic identity for a system with a single mechanical degree of freedom.

Equation (67) is the first instance in which the full logical chain of the paper—Newton's laws → internal energy → fiber bundle → work–heat decomposition → fiber averaging—produces a quantitative prediction that can be compared with experiment. Every step in the derivation is either a theorem of the framework or an explicit, named assumption (Assumption V.1, Assumption VI.3); no thermodynamic postulate has been invoked.

The Macroscopic First Law

The microscopic first law $\mathrm{d}U = \delta W + \delta Q$ holds pointwise on $\mathcal{E}$. We now show that fiber averaging produces a macroscopic identity on $\mathcal{B}$. The argument turns on a single asymmetry: U is constant on every fiber of the microcanonical ensemble, while the right-hand side is not.

Proposition VI.13 (Macroscopic First Law)**.** *Define the macroscopic heat form*

$$\delta\bar{Q} \;\coloneqq\; \int_{\mathcal{F}_\alpha} \delta Q\;\mathrm{d}\mu_\alpha^{\mathrm{mc}}. \tag{68}$$

Then

$$\mathrm{d}U \;=\; \delta\bar{W} + \delta\bar{Q}. \tag{69}$$

Proof. Integrate both sides of the microscopic first law against $\mathrm{d}\mu_\alpha^{\mathrm{mc}}$. By linearity:

$$\int_{\mathcal{F}_\alpha} \mathrm{d}U\;\mathrm{d}\mu_\alpha^{\mathrm{mc}} \;=\; \delta\bar{W} + \delta\bar{Q}. \tag{70}$$

It remains to evaluate the left-hand side. The measure μ_α^{mc} is supported on the energy shell $\Sigma_\alpha(U) = \{H_\alpha = U\}$. On this shell, $U = H_\alpha$ is identically equal to the fixed constant U—it does not vary across the fiber. The differential $\mathrm{d}U$, which records changes as α varies between thermodynamic states, is therefore constant on $\mathcal{F}_\alpha$ and passes through the integral:

$$\int_{\mathcal{F}_\alpha} \mathrm{d}U\;\mathrm{d}\mu_\alpha^{\mathrm{mc}} \;=\; \mathrm{d}U\cdot\underbrace{\int_{\mathcal{F}_\alpha}\mathrm{d}\mu_\alpha^{\mathrm{mc}}}_{=1} \;=\; \mathrm{d}U. \tag{71}$$

□

The proof exploits the fact that U is the same function whose constancy *defined* the fiber: every accessible microstate carries exactly the same energy, while the instantaneous generalized force $\mathcal{F}_j^{\mathrm{tot}}$ fluctuates across microstates. The two sides of the microscopic first law therefore behave differently under averaging—the left side passes through unchanged, the right side is genuinely averaged—and this asymmetry is what makes the macroscopic first law a consequence of the microscopic one rather than an independent postulate.

Remark VI.14 (Connection to the measured boundary force). The thermodynamic force $\bar{\mathcal{F}}_j^{\mathrm{tot}} = \langle\partial_{\alpha_j} H_\alpha\rangle_{\mathrm{mc}}$ is built from the conservative Hamiltonian alone (Proposition V.9). A pressure gauge, however, measures the total boundary force, including non-conservative impacts. For the two to agree, the microcanonical average of the non-conservative boundary power must vanish. This holds whenever the non-conservative forces satisfy a *dissipative parity condition*—oddness under momentum reversal $\mathbf{p}' \mapsto -\mathbf{p}'$—which is a defining property of all standard dissipative interactions (viscous drag, radiation reaction). The microcanonical measure is invariant under momentum reversal (since H_α is even in $\boldsymbol{\pi}$), so the average of any momentum-odd observable vanishes by symmetry. This ensures that $\bar{\mathcal{F}}_j^{\mathrm{tot}}$ coincides with the physically measured force at equilibrium.

Remark VI.15 (The canonical ensemble). In the canonical ensemble, H_α fluctuates across the fiber and the asymmetry exploited above is lost. The macroscopic internal energy must itself be defined as a fiber average:

$\bar{U}(\alpha,\beta) \coloneqq \langle H_\alpha\rangle_{\mu_\alpha^{\rm can}} = -\partial \ln Z/\partial\beta|_\alpha$. The macroscopic first law $\mathrm{d}\bar{U} = \delta\bar{W}_{\rm can} + \delta\bar{Q}_{\rm can}$ then requires accounting for the β-dependence of the measure itself.

VII. ENTROPY, THE INTEGRATING FACTOR, AND THE SECOND LAW

The macroscopic first law (Proposition VI.13) reads

$$\mathrm{d}U \;=\; \delta\bar{W} + \delta\bar{Q}, \qquad \delta\bar{Q} \;=\; \mathrm{d}U - \sum_j \bar{\mathcal{F}}_j^{\rm tot}\,\mathrm{d}\alpha_j. \quad (72)$$

The internal energy U is a state function; neither $\delta\bar{W}$ nor $\delta\bar{Q}$ is exact. We now show that $\delta\bar{Q}$ admits an integrating factor $1/T$, with $T\,\mathrm{d}S = \delta\bar{Q}$ for a state function $S = k_B\ln\Omega$. Neither S nor T is postulated; both emerge from a short computation once $\delta\bar{Q}$ is re-expressed in terms of the phase-space objects that define it.

The Macroscopic State Space

Definition VII.1 (Macroscopic state space)**.** The *macroscopic state space* is

$$\mathcal{M} \;\coloneqq\; \big\{(\alpha,U) : \alpha\in\mathcal{B},\; U > U_{\rm min}(\alpha)\big\}, \qquad (73)$$

where $U_{\rm min}(\alpha) \coloneqq \min_{\mathcal{F}_\alpha} H_\alpha$ is the ground-state energy. A *quasi-static process* is a smooth curve in $\mathcal{M}$.

Proposition VII.2 (Inexactness of $\delta\bar{Q}$)**.** *If $\bar{\mathcal{F}}_j^{\rm tot}(\alpha,U)$ depends non-trivially on U for some j, then $\delta\bar{Q}$ is not exact on $\mathcal{M}$.*

Proof. From (72), the coefficient of $\mathrm{d}U$ in $\delta\bar{Q}$ is 1. If $\delta\bar{Q} = \mathrm{d}f$, then $f = U + g(\alpha)$ for some g, so the coefficient of $\mathrm{d}\alpha_j$ in $\mathrm{d}f$ is independent of U. But the coefficient of $\mathrm{d}\alpha_j$ in $\delta\bar{Q}$ is $-\bar{\mathcal{F}}_j^{\rm tot}(\alpha,U)$, which depends on U by assumption. □

Re-expressing Macroscopic Heat in Terms of the Phase Volume

Define the *cumulative phase volume* and the *density of states*:

$$\begin{aligned}\Omega(\alpha,U) &\coloneqq \int_{\mathcal{F}_\alpha}\theta(U-H_\alpha)\,\mathrm{d}\Gamma_\alpha,\\ \Sigma(\alpha,U) &\coloneqq \frac{\partial\Omega}{\partial U} = \int_{\mathcal{F}_\alpha}\delta(H_\alpha-U)\,\mathrm{d}\Gamma_\alpha.\end{aligned} \qquad (74)$$

The thermodynamic force (Definition VI.12) is $\bar{\mathcal{F}}_j^{\rm tot} = \langle\partial_{\alpha_j}H_\alpha\rangle_{\rm mc}$ by construction. The key identity is the following.

Lemma VII.3 (Differentiation of the phase volume)**.** *Suppose the energy shell $\{H_\alpha = U\}$ lies in the interior of the configuration domain for all α. Then*

$$\left.\frac{\partial\Omega}{\partial\alpha_j}\right|_U \;=\; -\,\Sigma(\alpha,U)\,\bar{\mathcal{F}}_j^{\rm tot}(\alpha,U). \qquad (75)$$

Proof. Differentiate $\Sigma = \int\delta(H_\alpha-U)\,\mathrm{d}\Gamma_\alpha$ with respect to α_j at fixed U. The boundary assumption ensures the integrand has compact support, so differentiation under the integral is valid. By the chain rule for distributions:

$$\begin{aligned}\left.\frac{\partial\Sigma}{\partial\alpha_j}\right|_U &= -\frac{\partial}{\partial U}\int_{\mathcal{F}_\alpha}\frac{\partial H_\alpha}{\partial\alpha_j}\,\delta(H_\alpha-U)\,\mathrm{d}\Gamma_\alpha\\ &= -\frac{\partial}{\partial U}\big[\Sigma\,\bar{\mathcal{F}}_j^{\rm tot}\big].\end{aligned} \qquad (76)$$

Since $\Omega(\alpha,U) = \int_{U_{\rm min}}^U\Sigma(\alpha,U')\,\mathrm{d}U'$, differentiating with respect to α_j and substituting (76):

$$\left.\frac{\partial\Omega}{\partial\alpha_j}\right|_U \;=\; \int_{U_{\rm min}}^U\frac{\partial\Sigma}{\partial\alpha_j}\,\mathrm{d}U' \;=\; -\big[\Sigma\,\bar{\mathcal{F}}_j^{\rm tot}\big]_{U_{\rm min}}^U. \qquad (77)$$

At $U' = U_{\rm min}$, $\Sigma = 0$ (the energy shell at the ground state has zero measure), so the lower boundary vanishes. □

Assembling the full differential of Ω from $\partial_U\Omega = \Sigma$ and Lemma VII.3:

$$\begin{aligned}\mathrm{d}\Omega &= \Sigma\,\mathrm{d}U - \sum_j\Sigma\,\bar{\mathcal{F}}_j^{\rm tot}\,\mathrm{d}\alpha_j\\ &= \Sigma\left(\mathrm{d}U - \sum_j\bar{\mathcal{F}}_j^{\rm tot}\,\mathrm{d}\alpha_j\right)\\ &= \Sigma\,\delta\bar{Q}.\end{aligned} \qquad (78)$$

Therefore:

$$\delta\bar{Q} \;=\; \frac{\mathrm{d}\Omega}{\Sigma}. \qquad (79)$$

The heat form is the differential of the phase volume divided by the density of states. Both objects were already present in the microcanonical measure; equation (79) brings them to the surface.

The Integrating Factor

Given (79), finding an integrating factor is immediate. Multiplying by Σ/Ω:

$$\frac{\Sigma}{\Omega}\,\delta\bar{Q} \;=\; \frac{\mathrm{d}\Omega}{\Omega} \;=\; \mathrm{d}(\ln\Omega). \qquad (80)$$

Definition VII.4 (Entropy and absolute temperature)**.**

$$\begin{aligned}S(\alpha,U) &\coloneqq k_B\ln\Omega(\alpha,U),\\ \frac{1}{T(\alpha,U)} &\coloneqq \left.\frac{\partial S}{\partial U}\right|_\alpha = \frac{k_B\,\Sigma}{\Omega}.\end{aligned} \qquad (81)$$

Since $\Omega > 0$ and $\Sigma > 0$ on $\mathcal{M}$, $T > 0$ everywhere.

Theorem VII.5 (Entropy as integrating factor; Gibbs fundamental relation)**.** *$1/T$ is an integrating factor for $\delta\bar{Q}$:*

$$\mathrm{d}S \;=\; \frac{\delta\bar{Q}}{T}. \qquad (82)$$

Equivalently, substituting $\delta\bar{Q} = \mathrm{d}U - \delta\bar{W}$:

$$\mathrm{d}U \;=\; T\,\mathrm{d}S + \delta\bar{W}, \tag{83}$$

the Gibbs fundamental relation.

Proof. $\delta\bar{Q}/T = (k_B\Sigma/\Omega)\cdot(\mathrm{d}\Omega/\Sigma) = k_B\,\mathrm{d}\Omega/\Omega = \mathrm{d}(k_B\ln\Omega) = \mathrm{d}S$. ☐

Remark VII.6 (Why $S = k_B\ln\Omega$ and not $S = k_B\ln\Sigma$). The derivation rests on $\mathrm{d}\Omega = \Sigma\,\delta\bar{Q}$, which required the vanishing of the boundary term $\Sigma(\alpha, U_{\min}) = 0$ in Lemma VII.3. The shell entropy $S_\Sigma = k_B\ln\Sigma$ does not satisfy $\mathrm{d}S_\Sigma = \delta\bar{Q}/T_\Sigma$ exactly: the identity (76) introduces an additional term $\Sigma\,\partial_U\bar{\mathcal{F}}_j^{\mathrm{tot}}$ that vanishes only in the thermodynamic limit. The volume entropy $S = k_B\ln\Omega$ is exact at every N; this is the algebraic reason for its superiority, as recognised by Gibbs [21] and Hertz [26]. See [16, 27] for the modern discussion.

Remark VII.7 (Uniqueness of the integrating factor). If $\lambda\,\delta\bar{Q} = \mathrm{d}f$ for some smooth f, then since $\delta\bar{Q} = T\,\mathrm{d}S$, we have $\mathrm{d}f = \lambda T\,\mathrm{d}S$, whence $\mathrm{d}(\lambda T)\wedge\mathrm{d}S = 0$, so λT is a function of S alone. The freedom is a monotone reparametrisation of the entropy scale. The zeroth law (transitivity of thermal equilibrium) fixes this freedom to a single positive multiplicative constant—the choice of temperature unit.

Corollary VII.8 (Standard identities)**.** *We have the following identities:*

(i) *For a gas with $\alpha = V$: $\mathrm{d}U = T\,\mathrm{d}S - P\,\mathrm{d}V$, with $P = -\bar{\mathcal{F}}_V^{\mathrm{tot}}$.*
(ii) $\bar{\mathcal{F}}_j^{\mathrm{tot}} = -T\,(\partial S/\partial\alpha_j)|_U$.
(iii) *Maxwell relations:*

$$\partial_{\alpha_k}(\bar{\mathcal{F}}_j^{\mathrm{tot}}/T)|_U = \partial_{\alpha_j}(\bar{\mathcal{F}}_k^{\mathrm{tot}}/T)|_U\ .$$

The Second Law

Theorem VII.5 holds for quasi-static processes. For irreversible processes, the equality $\mathrm{d}S = \delta\bar{Q}/T$ is replaced by an inequality. The argument requires three ingredients.

a. Fine-grained entropy is conserved. Let ρ_t be the phase-space probability density at time t, with the fine-grained entropy $S_{\mathrm{fine}}(t) := -k_B\int\rho_t\ln\rho_t\,\mathrm{d}\Gamma_\alpha$. The Liouville equation $\partial_t\rho_t + \{\rho_t, H_\alpha\} = 0$ implies $\mathrm{d}S_{\mathrm{fine}}/\mathrm{d}t = 0$ for Hamiltonian dynamics: fine-grained entropy is exactly conserved.

b. Coarse-graining increases entropy. Partition $\mathcal{F}_\alpha$ into cells $\{\mathcal{C}_k\}$ and define the coarse-grained density $\bar{\rho}_t$ as the cell average of ρ_t. The coarse-grained entropy $S_{\mathrm{cg}} = -k_B\int\bar{\rho}_t\ln\bar{\rho}_t\,\mathrm{d}\Gamma_\alpha$ satisfies

$$S_{\mathrm{cg}}(t) \;\geq\; S_{\mathrm{fine}}(t) \tag{84}$$

by Jensen's inequality applied to the strictly convex function $x\ln x$, with equality iff ρ_t is constant on every cell.

c. Entropy non-decrease. Under the assumption of *molecular chaos*—that pre-collision states of interacting particles are statistically uncorrelated—the composite operation of Hamiltonian evolution followed by coarse-graining yields

$$S_{\mathrm{cg}}(t+\Delta t) \;\geq\; S_{\mathrm{fine}}(t+\Delta t) \;=\; S_{\mathrm{fine}}(t) \;=\; S_{\mathrm{cg}}(t), \tag{85}$$

where the inequality is (84), the first equality is Liouville conservation (Proposition omitted but standard), and the last equality uses the molecular chaos assumption: ρ_t is assumed to be effectively coarse-grained at each time step, so $S_{\mathrm{fine}}[\rho_t] = S_{\mathrm{cg}}[\bar{\rho}_t] = S_{\mathrm{cg}}(t)$. This is precisely where the Stosszahlansatz enters and where the arrow of time is inserted into the argument.

Remark VII.9 (Status of the molecular chaos assumption). The Stosszahlansatz [7] is the single irreversible input in the derivation. Loschmidt's reversibility objection [37] and Zermelo's recurrence objection [57] are both correct as mathematical theorems; the Stosszahlansatz asserts that the entropy-decreasing trajectories they identify are measure-theoretically negligible among all initial conditions compatible with the macroscopic state. For modern treatments, see [23, 35].

Theorem VII.10 (Clausius inequality)**.** *Let a system $\mathcal{S}$ undergo any process while in thermal contact with a reservoir at temperature T_R, with δQ the heat flowing into $\mathcal{S}$. Then*

$$\mathrm{d}S_{\mathcal{S}} \;\geq\; \frac{\delta Q}{T_R}, \tag{86}$$

with equality for quasi-static processes.

Proof. The composite $\mathcal{S}\cup\mathcal{R}$ is isolated, so $\mathrm{d}S_{\mathcal{S}} + \mathrm{d}S_{\mathcal{R}} \geq 0$ by (85). The reservoir responds quasi-statically by definition, so $\mathrm{d}S_{\mathcal{R}} = -\delta Q/T_R$ by Theorem VII.5. Substituting gives (86). ☐

Corollary VII.11 (The Second Law)**.** (i) Entropy increase. *For an isolated system ($\delta Q = 0$):* $\mathrm{d}S \geq 0$.
(ii) Direction of heat flow. *Two systems at $T_1 > T_2$ in thermal contact:* $\mathrm{d}S_{\mathrm{total}} = \delta Q(1/T_2 - 1/T_1) \geq 0$ *forces $\delta Q \geq 0$; heat flows from hot to cold.*

VIII. CONCLUSION

The paper has pursued a single question: can the decomposition of energy transfer into work and heat be derived from mechanics rather than postulated? The answer is affirmative, provided one formalises the physics of confinement as a fiber bundle.

The logical chain has six links. The work–energy theorem (Section II) and König's theorem (Section III) are standard results of Newtonian mechanics. The internal energy $U = K_{\mathrm{rel}} + \Phi$ (Section IV) is forced by the Poincaré lemma once one asks for the maximal state function absorbable from the microscopic work. The fiber bundle of confined configurations (Section V), equipped with a trivialization $\varphi_\alpha : \Omega_0 \to \Omega(\alpha)$, decomposes $\mathrm{d}U$ into horizontal (work) and vertical (heat) components by the chain rule—the microscopic first law $\mathrm{d}U = \delta W + \delta Q$ (Theorem V.7). Fiber averaging over the microcanonical measure (Section VI) produces the macroscopic first law $\mathrm{d}U = \delta\bar{W} + \delta\bar{Q}$ (Proposition VI.13), and re-expressing $\delta\bar{Q}$ in terms of the phase volume Ω yields entropy $S = k_B \ln \Omega$, temperature $1/T = k_B\Sigma/\Omega$, and the Gibbs fundamental relation $\mathrm{d}U = T\,\mathrm{d}S + \delta\bar{W}$ (Theorem VII.5). The second law follows from the Stosszahlansatz (Theorem VII.10).

Three assumptions beyond Newton's laws were required, each stated explicitly: the bundle regularity hypothesis (Assumption V.1), ergodicity of the Hamiltonian flow on the energy shell (Assumption VI.3), and molecular chaos (Remark VII.9). No thermodynamic postulate—no Clausius statement, no Kelvin–Planck statement, no Carathéodory axiom—was invoked.

Several features of the construction deserve comment.

The work–heat decomposition depends on the trivialization φ_α, which is determined by the physical apparatus implementing the boundary deformation. This is not an ambiguity but a reflection of the process-dependence of work and heat: different experimental protocols for achieving the same macroscopic change $\Delta\alpha$ produce different trivializations and hence different partitions of $\mathrm{d}U$. The total $\mathrm{d}U$ is invariant; the partition is not.

The connection induced by the trivialization is flat (Corollary V.6), so the path-dependence of the work form arises entirely from the fiber-dependence of the generalized force $\mathcal{F}_j^{\mathrm{tot}}$—the mechanical evolution of the microstate along different paths in $\mathcal{B}$—not from any geometric holonomy of the bundle.

The entropy $S = k_B \ln \Omega$ is the volume (Gibbs) entropy, not the surface (Boltzmann) entropy $k_B \ln \Sigma$. The algebraic reason is that the identity $\mathrm{d}\Omega = \Sigma\,\delta\bar{Q}$ (Lemma VII.3) holds exactly at every N, while the corresponding identity for Σ acquires corrections of order $1/N$. This provides a clean algebraic reason, within the present framework, for preferring the volume entropy over the surface entropy: $S = k_B \ln \Omega$ satisfies $\mathrm{d}S = \delta\bar{Q}/T$ exactly at every N, while $k_B \ln \Sigma$ does so only in the thermodynamic limit.

The framework has limitations that should be stated honestly. It applies to classical systems with a finite number of particles confined to bounded domains; quantum effects, long-range interactions requiring infinite-volume limits, and relativistic systems lie outside its scope. The ergodic hypothesis, while believed to hold for generic many-body systems, is rigorously established only for special classes (Sinai billiards [51], Anosov flows [3]); its failure for integrable systems means that the microcanonical measure is not uniquely selected for such systems within the present framework. The thermodynamic limit $N \to \infty$, which suppresses fluctuations and ensures the equivalence of ensembles [48, 52], is invoked informally (e.g., in the equipartition theorem and in the derivation of the canonical ensemble) but not treated with the rigour of mathematical statistical mechanics. The $N!$ and h^{3N} factors in the grand canonical ensemble are quantum in origin (Remark VI.10) and cannot be derived classically.

Two directions for further development are natural. First, the quantum-mechanical analogue of the construction—replacing the classical configuration-space bundle $\mathcal{E}_\mathcal{Q}$ with the bundle of Hilbert spaces $\{L^2(\Omega(\alpha))\}_\alpha$ and the Ehresmann connection with the Berry connection [6, 50]—would provide a unified geometric framework for classical and quantum thermodynamics, and would make contact with the gauge-theoretic approaches of Roberts [46], Céleri and collaborators [12, 19, 43]. Second, extending the macroscopic first law to the canonical ensemble—where U fluctuates across the fiber and the asymmetry exploited in Proposition VI.13 is absent—requires a different averaging argument and leads naturally to the free energy $F = U - TS$ as the appropriate thermodynamic potential; this is left for future work.

The central message is modest. The work–heat decomposition is geometry: it is the horizontal–vertical splitting of an exact 1-form on a fiber bundle, induced by the physics of confinement. Entropy is algebra: it is the integrating factor that falls out when the macroscopic heat form is written in terms of the phase volume. The second law is statistics: it is the non-decrease of coarse-grained entropy under the molecular chaos hypothesis. Each of these statements has a precise mathematical content, and none requires a thermodynamic axiom.

[1] Ralph Abraham and Jerrold E. Marsden. *Foundations of Mechanics*. Benjamin/Cummings, Reading, MA, 2nd edition, 1978.

[2] Yakir Aharonov and David Bohm. Significance of electromagnetic potentials in the quantum theory. *Phys. Rev.*, 115:485–491, 1959.

[3] Dmitri V. Anosov. Geodesic flows on closed Riemannian manifolds of negative curvature. *Proc. Steklov Inst. Math.*, 90:1–235, 1967. Translated from Russian; establishes ergodicity of geodesic flows on negatively curved manifolds.

[4] Vladimir I. Arnold. *Mathematical Methods of Classical Mechanics*. Springer, New York, 2nd edition, 1989. Translated by K. Vogtmann and A. Weinstein.

[5] Roger Balian and Paul Valentin. Hamiltonian structure of thermodynamics with gauge. *Eur. Phys. J. B*, 21:269–282, 2001.

[6] Michael V. Berry. Quantal phase factors accompanying adiabatic changes. *Proc. R. Soc. Lond. A*, 392:45–57, 1984.

[7] Ludwig Boltzmann. Weitere Studien über das Wärmegleichgewicht unter Gasmolekülen. *Sitzungsber. Akad. Wiss. Wien*, 66:275–370, 1872. Contains the H-theorem and the molecular chaos assumption.

[8] Raoul Bott and Loring W. Tu. *Differential Forms in Algebraic Topology.* Springer, New York, 1982.

[9] Alessandro Bravetti, Hans Cruz, and Diego Tapias. Contact Hamiltonian mechanics. *Ann. Phys.*, 376:17–39, 2017.

[10] Herbert B. Callen. *Thermodynamics and an Introduction to Thermostatistics.* Wiley, New York, 2nd edition, 1985.

[11] Constantin Carathéodory. Über die Bestimmung der Energie und der absoluten Temperatur mit Hilfe von reversiblen Prozessen. *Math. Ann.*, 67:355–386, 1909.

[12] Lucas C. Céleri and Łukasz Rudnicki. Gauge-invariant quantum thermodynamics: consequences for the first law. *arXiv preprint*, 2024.

[13] Rudolf Clausius. Über einen auf die Wärme anwendbaren mechanischen Satz. *Ann. Phys.*, 217:124–130, 1870. Original statement and proof of the virial theorem.

[14] Earl A. Coddington and Norman Levinson. *Theory of Ordinary Differential Equations.* McGraw-Hill, New York, 1955.

[15] Paul A. M. Dirac. Quantised singularities in the electromagnetic field. *Proc. R. Soc. Lond. A*, 133:60–72, 1931.

[16] Jörn Dunkel and Stefan Hilbert. Consistent thermostatistics forbids negative absolute temperatures. *Nature Phys.*, 10:67–72, 2014.

[17] Charles Ehresmann. Les connexions infinitésimales dans un espace fibré différentiable. *Colloque de Topologie (Bruxelles, 1950)*, pages 29–55, 1950.

[18] Enrico Fermi. *Thermodynamics.* Dover, New York, 1956. Reprint of the 1936 edition.

[19] Lorenzo Ferrari, Łukasz Rudnicki, and Lucas C. Céleri. Gauge theory of quantum thermodynamics. *arXiv preprint*, 2025.

[20] François Gay-Balmaz and Hiroaki Yoshimura. From Lagrangian mechanics to nonequilibrium thermodynamics: a variational perspective. *Entropy*, 21(1):8, 2019.

[21] Josiah Willard Gibbs. *Elementary Principles in Statistical Mechanics.* Charles Scribner's Sons, New York, 1902.

[22] Herbert Goldstein, Charles P. Poole, and John L. Safko. *Classical Mechanics.* Addison-Wesley, San Francisco, 3rd edition, 2002.

[23] Sheldon Goldstein. Boltzmann's approach to statistical mechanics. *Chance in Physics: Foundations and Perspectives*, pages 39–54, 2001.

[24] Jean-Pierre Hansen and Ian R. McDonald. *Theory of Simple Liquids.* Academic Press, Oxford, 4th edition, 2013.

[25] Robert Hermann. *Geometry, Physics, and Systems.* Marcel Dekker, New York, 1973. First systematic use of contact geometry in thermodynamics.

[26] Paul Hertz. Über die mechanischen Grundlagen der Thermodynamik. *Ann. Phys.*, 338:225–274, 537–552, 1910. Systematic justification of the volume entropy $S = k_B \ln \Omega$.

[27] Stefan Hilbert, Peter Hänggi, and Jörn Dunkel. Thermodynamic laws in isolated systems. *Phys. Rev. E*, 90:062116, 2014.

[28] Sergio A. Hojman and Lawrence C. Shepley. No Lagrangian? No quantization! *J. Math. Phys.*, 32:142–146, 1991. Discusses mass-tensor ambiguities in inverse problems of mechanics.

[29] Kerson Huang. *Statistical Mechanics.* Wiley, New York, 2nd edition, 1987.

[30] Edwin T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, 1957. Maximum-entropy derivation of the microcanonical and canonical measures.

[31] Aleksandr Ya. Khinchin. *Mathematical Foundations of Statistical Mechanics.* Dover, New York, 1949. Translated by G. Gamow; contains a rigorous proof of the equipartition theorem.

[32] Shoshichi Kobayashi and Katsumi Nomizu. *Foundations of Differential Geometry*, volume I. Wiley, New York, 1963.

[33] Samuel König. De universali principio aequilibrii et motus. *Nova Acta Eruditorum*, pages 125–135, 162–176, 1751. The original statement of the kinetic-energy decomposition theorem.

[34] Valery V. Kozlov. Integrability and non-integrability in Hamiltonian mechanics. *Russian Math. Surveys*, 38(1):1–76, 1983.

[35] Joel L. Lebowitz. Statistical mechanics: a selective review of two central issues. *Rev. Mod. Phys.*, 71:S346–S357, 1999.

[36] Elliott H. Lieb and Jakob Yngvason. The physics and mathematics of the second law of thermodynamics. *Phys. Rep.*, 310:1–96, 1999.

[37] Josef Loschmidt. Über den Zustand des Wärmegleichgewichtes eines Systems von Körpern mit Rücksicht auf die Schwerkraft. *Sitzungsber. Akad. Wiss. Wien*, 73:128–142, 1876.

[38] Lawrence Markus and Kenneth R. Meyer. Generic Hamiltonian dynamical systems are neither integrable nor ergodic. *Mem. Amer. Math. Soc.*, 144:1–52, 1974.

[39] Jerrold E. Marsden and Tudor S. Ratiu. *Introduction to Mechanics and Symmetry.* Springer, New York, 2nd edition, 1999.

[40] Ryszard Mrugała. Geometrical formulation of equilibrium phenomenological thermodynamics. *Rep. Math. Phys.*, 14:419–427, 1978.

[41] Ryszard Mrugała. Geometrical methods in thermodynamics. *Thermodynamics of Energy Conversion and Transport*, pages 257–285, 2000.

[42] Mikio Nakahara. *Geometry, Topology and Physics.* Taylor & Francis, Boca Raton, 2nd edition, 2003.

[43] Daniel Pernambuco-Wise and Lucas C. Céleri. Principal fiber bundles in the gauge theory of quantum thermodynamics. *arXiv preprint*, 2025.

[44] Alfred B. Pippard. *Elements of Classical Thermodynamics.* Cambridge University Press, Cambridge, 1957.

[45] Henri Poincaré. Analysis situs. *J. École Polytech.*, 1:1–121, 1895.

[46] Nicholas Roberts. Heat as a gauge connection: a fiber bundle formulation of thermodynamics. *arXiv preprint*, 2025.

[47] Walter Rudin. *Real and Complex Analysis.* McGraw-Hill, New York, 3rd edition, 1987.

[48] David Ruelle. *Statistical Mechanics: Rigorous Results.* World Scientific, Singapore, 1999. Reprint of the 1969 edition.

[49] George Ruppeiner. Thermodynamics: a Riemannian geometric model. *Phys. Rev. A*, 20:1608–1613, 1979.

[50] Barry Simon. Holonomy, the quantum adiabatic theorem, and Berry's phase. *Phys. Rev. Lett.*, 51:2167–2170, 1983.

[51] Yakov G. Sinai. Dynamical systems with elastic reflections: ergodic properties of dispersing billiards. *Russian Math. Surveys*, 25(2):137–189, 1970.

[52] Hugo Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478:1–69, 2009.

[53] Arjan van der Schaft and Bernhard Maschke. Geometry of thermodynamic processes. *Entropy*, 20(12):925, 2018.

Extended version: *ibid.* 23 (2021) 1604.

[54] Frank W. Warner. *Foundations of Differentiable Manifolds and Lie Groups.* Springer, New York, 1983.

[55] Frank Weinhold. Metric geometry of equilibrium thermodynamics. *J. Chem. Phys.*, 63:2479–2483, 1975.

[56] Tai Tsun Wu and Chen Ning Yang. Concept of nonintegrable phase factors and global formulation of gauge fields. *Phys. Rev. D*, 12:3845–3857, 1975.

[57] Ernst Zermelo. Über einen Satz der Dynamik und die mechanische Wärmetheorie. *Ann. Phys.*, 293:485–494, 1896.