

A Finite-State Gibbs Construction from a Recognition Cost

M Simons¹ and J Washburn¹

¹Recognition Science, Recognition Physics Institute, Austin, Texas, USA

E-mail: msimons@recognitionphysics.org

Abstract. On a finite outcome space, the canonical Gibbs distribution is usually obtained by maximizing Shannon entropy at fixed mean of an externally supplied energy functional. This paper studies the finite-state consequences of a ratio-cost construction instead: after adopting the normalized d'Alembert degree-two closure called the Recognition Composition Law (RCL), with unit log-curvature calibration at the reference ratio, the continuous nontrivial positive branch is $J(x) = \frac{1}{2}(x+x^{-1}) - 1 = \cosh(\log x) - 1$. Given the induced cost vector $X_\omega = J(r_\omega)$, multinomial counting and convex duality recover the finite-state Gibbs weights and the identity $F_R(q) - F_R(p) = T_R D_{\text{KL}}(q||p)$; the entropy-maximization steps are classical once the cost is fixed. New technical content includes a non-asymptotic Stirling bound and soft-shell constrained-type theorems for real-valued costs. A three-state example compares the Gibbs law to squared-log, affinity-as-energy, and Tsallis alternatives at the same cost vector and mean-cost constraint, with sample-size power calculations at fixed RCL ground truth. The framework is conditional on axioms (A1)–(A3) and restricted to finite outcome spaces with strictly positive weights; it does not derive the composition law from a more primitive principle.

AMS classification scheme numbers: 82B03, 82B05, 60F10, 94A15

Keywords: classical statistical mechanics, equilibrium and non-equilibrium; entropy; information theory; large deviations in non-equilibrium systems; rigorous results in statistical mechanics.

1. Introduction

On a finite outcome space, adopting the normalized d’Alembert degree-two closure called the Recognition Composition Law (RCL) of Section 1.2, together with unit log-curvature at the reference ratio, selects an energy-like cost on the dimensionless ratios that label outcomes. Once this cost is fixed, the standard finite-state machinery of multinomial counting and convex duality gives the Boltzmann–Gibbs measure, the Helmholtz free energy, and the Kullback–Leibler relaxation theorem. The axiomatic commitment is not derived from a more primitive principle; the paper studies its statistical-mechanical consequences.

1.1. Motivation and scope

Statistical mechanics rests on the Boltzmann–Gibbs distribution [1, 2]

$$p_i = \frac{e^{-E_i/k_B T}}{Z}, \quad Z = \sum_i e^{-E_i/k_B T}, \quad (1)$$

where E_i are energy levels, T is temperature, and k_B is Boltzmann’s constant. The canonical form (1) is not a postulate: it is derived from more basic postulates—typically the equal *a priori* probability hypothesis, or the microcanonical ensemble in the thermodynamic limit—and the derivation takes the energy functional E_i as given. Rigorous mathematical foundations for this derivation exist [3]. Jaynes’s maximum-entropy program [4, 5, 6] reframes the Gibbs distribution as the result of maximizing Shannon entropy subject to an energy constraint, but the energy functional is still taken as given. The Shore–Johnson axiomatization [7, 8] fixes Shannon entropy as the unique self-consistent inference rule under linear constraints, but again leaves the choice of constraint functional open. Generalizations such as Tsallis non-extensive statistics [9] and Naudts’s ϕ -exponential families [10] replace the entropy functional itself; for a modern survey of the large-deviation framework that organizes much of this material, see [11].

On a finite configuration space, the canonical construction can be phrased entirely in terms of constrained empirical compositions: a type plays the role of a macrostate, its multinomial weight determines an entropy density, and the Gibbs measure appears as the maximizer at fixed mean cost. This is the variational picture that underlies equivalence of ensembles and large-deviation rate functions in classical lattice and mean-field models. Here it is specialized to a single affine cost shell on a finite, strictly positive probability space, with a cost J tied to ratio fluctuations rather than to a spatial Hamiltonian.

What is fixed and what is left as input. We work on a finite outcome space throughout, and the framework is restricted to that setting; continuum configuration spaces, hard-core

or singular interactions, quantum reduced states, spatially extended Gibbs measures, ensemble equivalence in the sense of Ruelle [3] or Lanford [12], and the Boltzmann H -theorem for many-particle systems lie outside its scope. Within that setting, the framework supplies a single notational thread $X_\omega = J(r_\omega)$ connecting (a) a multiplicative–reciprocal functional equation on positive ratios, (b) the resulting closed-form cost, (c) the multinomial counting that yields Shannon entropy as the rate function, and (d) the canonical ensemble on the corresponding affine cost shell. The paper’s nonstandard input is the cost-determining functional equation and its finite-state consequences; the remaining variational and large-deviation steps are the standard backbone shared with Jaynes’s program. Modeling freedom is thereby relocated from the choice of cost functional to the choice of input ratios and a single calibration constant, not eliminated: specifying $\{r_\omega\}$ is itself a modeling commitment, and Section 8.5 shows that the choice of the d’Alembert composition law over its Cauchy, Tsallis, or Naudts alternatives is a further modeling commitment with empirically distinguishable consequences (Section 7.5).

Section 1.3 states the contribution and its expository and methodological components explicitly; Table 1 provides a roadmap of where each formal result enters the chain.

What is new here. In this sense, the novelty is not a new entropy principle or a new proof of the Gibbs variational principle, but the relocation of the modeling choice to dimensionless ratio data and a reciprocal composition law for their scalar cost. Section 1.3 and Table 1 record the technical components that follow once that commitment is adopted.

1.2. Recognition Composition Law and fixed cost

Throughout the paper we abbreviate “Recognition Composition Law” as RCL: it is the multiplicative–reciprocal functional equation (2) below, which prescribes how the scalar cost J of a positive dimensionless ratio x combines with that of a second ratio y when the two are merged (xy) and compared (x/y).

1.2.1. Physical setting: ratios as primitive bookkeeping variables. The RCL arises in a ratio-cost framework [13] in which positive dimensionless ratios, rather than energy levels, serve as the primitive bookkeeping variable for the outcomes of a finite system. The natural physical instances are reversible Markov chains and chemical reaction networks: each assigns a positive dimensionless ratio to every microscopic process or state, and each composes those ratios multiplicatively rather than additively.

Two concrete instances illustrate the point. In a reversible continuous-time Markov chain on a state space Ω with stationary distribution π and transition rates W_{ij} , attach to each state $a \in \Omega$ a dimensionless ratio $r_a := \pi_a/\pi_{a_0}$ relative to a chosen reference state a_0 (so $r_{a_0} = 1$). Detailed balance $\pi_a W_{ab} = \pi_b W_{ba}$ then assigns the edge ratio $r_b/r_a = \pi_b/\pi_a = W_{ab}/W_{ba}$ to the edge $a \leftrightarrow b$. When two such edges are concatenated, the corresponding odds ratios multiply: the product $r_b/r_a \cdot r_c/r_b = r_c/r_a$ is the natural composition. In a reaction network, equilibrium constants $K_{\text{eq}} = [B]/[A]$ for $A \rightleftharpoons B$ play

the same role—combining two reactions $A \rightleftharpoons B$ and $B \rightleftharpoons C$ multiplies their equilibrium constants because affinities $\Delta\mu/k_{\text{B}}T = \log K_{\text{eq}}$ add additively while the dimensionless ratios themselves multiply. As a reminder, the structural features (S1)–(S3) below are modeling assumptions imported from the bookkeeping conventions used alongside detailed balance and chemical equilibrium, and they are not implied by detailed balance itself.

Modeling conditions on a ratio-cost. We adopt three structural features of detailed-balance and equilibrium-constant bookkeeping as modeling assumptions on the scalar cost $J : \mathbb{R}_{>0} \rightarrow \mathbb{R}$; they are not implied by detailed balance or any more primitive physical principle, but summarize the symmetries we import from that bookkeeping context:

- (S1) Reciprocity (orientation invariance). The same edge $a \leftrightarrow b$ can be read in either direction, so the ratios r and r^{-1} describe the same reversible process; we require $J(x) = J(x^{-1})$.
- (S2) Reference vanishing. The reference ratio $x = 1$ corresponds to balanced odds, so $J(1) = 0$.
- (S3) Closure under products and quotients. Given two positive ratios x, y , the algebraic operations on $\mathbb{R}_{>0}$ produce two further positive ratios—the product xy (path composition) and the quotient x/y (relative odds of two parallel edges). We require that the joint cost $J(xy) + J(x/y)$ be a function only of $J(x)$ and $J(y)$, symmetric under $(x, y) \mapsto (y, x)$ and $(x, y) \mapsto (x^{-1}, y^{-1})$.

Conditions (S1)–(S3) do not determine a unique J : many functions satisfy them, and (S3) only names a class of composition rules without selecting one. The d’Alembert law (2) below is a separate postulate.

Choice of composition rule. We single out the d’Alembert composition law

$$J(xy) + J(x/y) = 2J(x)J(y) + 2J(x) + 2J(y),$$

i.e. equation (2) below, by adding one further postulate: $J(xy) + J(x/y)$ depends polynomially on $J(x)$ and $J(y)$ with total degree at most two, and the only constant term is the one forced by (S2). Symmetry of $J(xy) + J(x/y)$ under $(x, y) \mapsto (y, x)$ is automatic from (S1) and is not a separate postulate. This degree-two closure is a substantive modeling commitment, not a consequence of (S1)–(S3) or of detailed balance alone. The motivation for adopting it is twofold. First, if the composition law is required to be jointly analytic in $J(x)$ and $J(y)$ near the reference state $x = y = 1$, the constant, linear, and bilinear terms collected in (2) are exactly the terms that survive after imposing (S1)–(S2). Degree-two closure thus selects the simplest analytic truncation that is compatible with (S1)–(S3) and admits a non-constant solution with $J(1) = 0$. Second, the resulting two-variable polynomial dependence is observable: the joint cost $J(xy) + J(x/y)$ depends only on the individual costs $J(x), J(y)$ and their product $J(x)J(y)$, not on any further function of the underlying ratios.

Why degree two rather than higher order. The degree-two closure is the minimal analytic closure that includes a genuine pairwise interaction between two ratio-cost observations. A purely additive closure in $J(x)$ and $J(y)$ gives the squared-log branch in log-coordinates; adding the bilinear term $J(x)J(y)$ is the first non-additive correction compatible with a two-ratio composition rule. Operationally, this means that the combined cost of the two observable ratios xy and x/y is determined by $J(x)$, $J(y)$, and their pairwise product, with no auxiliary state, memory variable, or higher-order coupling. The closure is also testable as a finite relation among four observable costs, $J(xy)$, $J(x/y)$, $J(x)$, and $J(y)$. Higher-degree or non-polynomial closures are possible and are not ruled out by (S1)–(S3); the present paper treats degree two as the minimal non-additive analytic closure and studies the consequences of that disciplined modeling choice. A generic degree-three closure compatible with (S1), (S2), $J(1) = 0$, and a non-constant continuous solution is tightly constrained, and no closed-form analogue of cosh is known for the cubic case; that is the operational reason the manuscript restricts to degree two rather than opening a separate classification project.

Among the natural low-degree alternatives considered here are the trivial branch $J \equiv 0$, the additive d’Alembert/Cauchy branch in log-coordinates $G(u+v) + G(u-v) = 2G(u) + 2G(v)$ (whose continuous solutions $J = c(\log x)^2$ recover the squared-log surrogate revisited in Section 2), and the multiplicative Cauchy law $J(xy) = J(x)J(y)$ (incompatible with $J(1) = 0$ unless $J \equiv 0$). The list is illustrative rather than exhaustive once the d’Alembert postulate is relaxed; Section 8.5 collects a fuller tabulation of alternatives and their maximum-entropy distributions.

In log-coordinates the closure has a transparent form: writing $G(u) := J(e^u)$, the d’Alembert composition becomes the d’Alembert functional equation (6), and Theorem 1.1 below selects $G(u) = \cosh u - 1$. The squared-log surrogate $G(u) = u^2/2$ corresponds to the additive Cauchy law and agrees with the d’Alembert solution to second order at $u = 0$; Lemma 2.1 shows that the d’Alembert solution dominates pointwise everywhere else.

1.2.2. The composition law and its $y = 1$ reduction. The starting point is the d’Alembert-type composition law on $\mathbb{R}_{>0}$ used by Washburn and Zlatanović [13] (their equation (3)); multiplicative functional equations of this family are classical [14, 15]. The degree-two condition by itself leaves the normalization of the bilinear interaction as a modeling choice; in this paper the RCL is the normalized d’Alembert closure obtained by fixing that coefficient as in equation (2) below, after which the unit log-curvature calibration fixes the remaining scale of the continuous solution.

$$J(xy) + J(x/y) = 2J(x)J(y) + 2J(x) + 2J(y), \quad (2)$$

for an unknown cost $J : \mathbb{R}_{>0} \rightarrow \mathbb{R}$. The left-hand side aggregates the cost of the merged ratio xy with the cost of the compared ratio x/y ; the right-hand side combines the individual costs $J(x)$ and $J(y)$ with their multiplicative interaction $J(x)J(y)$. We take (2) as the composition rule adopted in the Recognition framework [13] and study

its calibrated consequences on a finite outcome space. An extension of this framework to additional cost structures appears in [16].

Setting $y = 1$ in (2) gives

$$2J(x) = 2J(x)J(1) + 2J(x) + 2J(1).$$

Subtracting $2J(x)$ from both sides leaves $0 = 2J(x)J(1) + 2J(1)$, and factoring $2J(1)$ on the right yields

$$0 = 2J(x)J(1) + 2J(1) = 2J(1)(J(x) + 1), \quad (3)$$

valid for every $x > 0$. Thus either $J(1) = 0$ or $J(x) = -1$ for every $x > 0$ (hence $J \equiv -1$) [13]; in the former case, reciprocity $J(x) = J(x^{-1})$ follows by setting $x = 1$ in (2) and using $J(1) = 0$ [13]. On the nontrivial branch selected below, $G(u) = J(e^u)$ satisfies $G(u) = \cosh u - 1$ near $u = 0$, hence $G'(0) = 0$ and the cost vanishes to first order at the reference ratio.

1.2.3. Calibration and classification. The composition law (2) together with $J(1) = 0$ admits the one-parameter family $x \mapsto \cosh(\lambda \log x) - 1$ of admissible solutions, indexed by the strength $\lambda \geq 0$ of the cost penalty. To fix this remaining scale, we calibrate J at the reference state by its second derivative in log-coordinates. Define the log-curvature of J at the reference state by

$$\kappa(J) := \lim_{t \rightarrow 0} \frac{2J(e^t)}{t^2}$$

whenever the limit exists; equivalently $\kappa(J) = \lim_{x \rightarrow 1} 2J(x)/(\log x)^2$. The condition $\kappa(J) = 1$ excludes the spurious solution $J \equiv -1$ and fixes the scale in the one-parameter family $x \mapsto \cosh(\lambda \log x) - 1$ of solutions of (2) with $J(1) = 0$ selected by Theorem 1.1. The second-order Taylor coefficient is the natural calibration handle because (S1)–(S2) force the first-order coefficient to vanish. The choice $\kappa(J) = 1$ is the natural normalization: it is the unique scale at which $D_J(q||p)$ reduces, near $q = p$, to the standard half-Neyman χ^2 statistic $\frac{1}{2}\chi^2(q||p)$ on the simplex (Proposition 3.1 (ii)), and equivalently the unique scale at which the squared-log surrogate $\frac{1}{2}(\log x)^2$ from Lemma 2.1 agrees with J to second order at $x = 1$. Any other choice $\kappa(J) = \lambda^2$ would amount to an overall rescaling of β and T_R in Sections 5–6 without altering the Gibbs measure or the free-energy–Kullback–Leibler identity.

Theorem 1.1 below is the central algebraic input to the rest of the paper. It matches the classification proved in the formal foundations of Washburn and Zlatanović [13] under measurability and a local integrability hypothesis on logarithmic coordinates.

Theorem 1.1 (Classification of measurable solutions of the RCL; selection by calibration). *Let $J : \mathbb{R}_{>0} \rightarrow \mathbb{R}$ be Lebesgue measurable and satisfy the composition law (2). Define $G(u) := J(e^u)$ for $u \in \mathbb{R}$, and assume the local integrability condition*

$$\int_{-\varepsilon}^{\varepsilon} |G(u)| du < \infty \quad \text{for some } \varepsilon > 0, \quad (4)$$

as in Washburn and Zlatanović [13]. Then either $J \equiv -1$, or $J(1) = 0$ and there exists $\lambda \geq 0$ with

$$J(x) = \cosh(\lambda \log x) - 1 \quad \text{or} \quad J(x) = \cos(\mu \log x) - 1 \quad (\mu \geq 0).$$

Imposing in addition the calibration $\kappa(J) = 1$ excludes the spurious branch $J \equiv -1$, the cosine branch (which has $\kappa = -\mu^2 \leq 0$), and the trivial $\lambda = 0$ representative; the remaining solution is $\lambda = 1$, giving

$$J(x) = \frac{1}{2}(x + x^{-1}) - 1. \quad (5)$$

Proof. First set $G(u) := J(e^u)$. With $x = e^u, y = e^v$, equation (2) becomes

$$G(u+v) + G(u-v) = 2G(u)G(v) + 2G(u) + 2G(v), \quad u, v \in \mathbb{R}, \quad (6)$$

and G is measurable because J is. Setting $u = v = 0$ gives $G(0)^2 + G(0) = 0$, so $G(0) \in \{0, -1\}$. If $G(0) = -1$, then (3) forces $G(u) \equiv -1$, and $\kappa(J) = \lim_{u \rightarrow 0} -2/u^2 = -\infty$, contradicting the calibration $\kappa(J) = 1$. Hence $G(0) = 0$.

Now set $H(u) := G(u) + 1$; then $H(0) = 1$ and substituting $G = H - 1$ into (6) gives

$$H(u+v) + H(u-v) = 2H(u)H(v), \quad u, v \in \mathbb{R}, \quad (7)$$

the classical d'Alembert equation [14, Ch. 3, §2]. Under measurability of H (inherited from J) together with (4), Washburn and Zlatanović [13] show that every solution of (7) with $H(0) = 1$ coincides pointwise with a continuous solution; in particular H takes one of the classical forms

$$H(u) = \cos(\mu u) \quad \text{or} \quad H(u) = \cosh(\lambda u), \quad \mu, \lambda \geq 0, \quad (8)$$

as in the continuous classification of Aczél [14, Theorem 5, p. 177] and Kuczma–Choczewski–Ger [15, §13.1]. (Without such a regularity hypothesis, pathological non-measurable solutions of (2) exist [13].)

The calibration selects the hyperbolic branch. For the cosine branch, $\kappa(J) = \lim_{u \rightarrow 0} 2(\cos \mu u - 1)/u^2 = -\mu^2 \leq 0$, which cannot equal 1. For the hyperbolic branch,

$$1 = \kappa(J) = \lim_{u \rightarrow 0} \frac{2(\cosh \lambda u - 1)}{u^2} = \lambda^2,$$

so $\lambda = 1$, $G(u) = \cosh u - 1$, and $J(e^u) = \cosh u - 1$. Substituting $u = \log x$ and using $\cosh(\log x) = \frac{1}{2}(x + x^{-1})$ gives (5). \square

The measurability and local integrability hypotheses align the present statement with the Washburn–Zlatanović formal treatment [13]; they also exclude the non-measurable pathologies recalled there. Continuity is not needed separately: the selected branch is smooth on $\mathbb{R}_{>0}$ once (5) is identified.

The function J measures the cost of a ratio x deviating from unity. It satisfies $J(x) \geq 0$ with equality if and only if $x = 1$, and $J(x) = J(1/x)$.

1.3. Assumptions, contributions, and summary of results

The primitive modeling assumptions are:

- (A1) The composition law (2) for $J : \mathbb{R}_{>0} \rightarrow \mathbb{R}$, with motivation discussed in Section 1.2.
- (A2) Unit log-curvature: $\kappa(J) = 1$.
- (A3) Lebesgue measurability of J on $\mathbb{R}_{>0}$ together with the local integrability condition (4) on $G(u) = J(e^u)$.

Theorem 1.1 shows that, conditional on (A1)–(A3), the selected cost is $J(x) = \frac{1}{2}(x + x^{-1}) - 1$. The integrability condition in (4) is the same standing hypothesis used in Washburn and Zlatanović [13] to rule out pathological non-measurable solutions while retaining measurability in place of continuity.

What is and what is not derived. The axioms (A1)–(A3) are not derived from a more primitive principle. (A1) is the d’Alembert composition (2); we further restrict to it from the broader (S1)–(S3) class by adopting an analytic-degree-2 closure (Section 1.2). (A2) fixes a normalization. (A3) is a measurability and local-integrability regularity package matching [13]. Joint adoption of these gives the closed form $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ and from there the standard exponential-family / Gibbs structure. The framework is conditional on the modeling commitment (A1)–(A3); it does not derive that commitment, and Section 8.5 catalogues the alternatives that become admissible if any one of (S1)–(S3) or the degree-2 closure is dropped.

Contributions and expository components. The construction keeps the cost-selection, counting, and variational steps explicit. Its components are:

- (N1, expository) The cost-classification theorem under (A1)–(A3) (Theorem 1.1), proved in the measurability plus local integrability regime of Washburn and Zlatanović [13]. The proof reduces (2) to the d’Alembert functional equation in log-coordinates [14, 15] and invokes their classification of measurable solutions before applying the calibration $\kappa(J) = 1$.
- (N2, methodological) We record an explicit non-asymptotic Stirling bound for the multinomial weights (Proposition 4.1), with all constants tracked through Robbins’ two-sided Stirling inequalities [17]. This is a direct application of a 1955 inequality, not a new estimate; its purpose is to make the regime of validity of the leading-order rate equality of Theorem 4.4 (i.e. $N \gg K \log N$) checkable at finite N, K .
- (N2-soft, methodological) The soft-shell constrained multinomial theorem (Theorem 4.4) for real-valued costs and shrinking tolerance windows, together with its multi-constraint extension (Proposition 8.1), removes the exact rational-cost restriction needed

for the exact-shell corollary (Corollary 4.5) in the experimentally relevant setting in which measured ratios and costs are not rationally engineered.

(N3, expository)

We unify the four steps under a common notation $\{r_\omega\} \rightarrow J \rightarrow X \rightarrow (\text{Gibbs}, F_R, D_{\text{KL}})$ and add (i) a tabulation of alternative composition laws and their maximum-entropy distributions (Section 8.5, Table 8) and (ii) a side-by-side numerical comparison with Tsallis q_T -exponentials at the same cost vector and constraint (Section 7.5). The alternative-composition tabulation lets a reader judge which parts of the cost selection come from (S1)–(S3) alone and which parts come from the additional degree-two closure.

The practical role of (N3) is inferential rather than thermodynamic: when the primary microscopic data are dimensionless ratios—odds ratios in a reversible Markov chain estimated from trajectories, equilibrium constants in a chemical network measured spectroscopically, detailed-balance ratios in a single-molecule experiment—the present chain converts those ratios directly into a finite-state canonical ensemble without an intervening Hamiltonian assignment. Modeling freedom is relocated from the choice of energy functional to the choice of ratio assignment together with the d’Alembert composition law (Section 1.2); whether that relocation is useful depends on whether ratios or energies are the more directly accessible objects in the application at hand. Sections 4.5 and 8.4 discuss representative settings, Section 7.6 quantifies the experimental sensitivity required to distinguish the RCL prediction from naive log-cost (“affinity-as-energy”) and Tsallis alternatives, and Section 8.3 relates the RCL cost $J(r)$ to the antisymmetric affinity $\log r$ of stochastic thermodynamics on the same edge data.

After J is fixed, the remaining statistical-mechanical machinery is standard: the method-of-types and large-deviation steps (Theorem 4.4), the finite-state Gibbs variational principle (Proposition 5.3), the free-energy–KL identity (Proposition 6.1 and its corollaries), and detailed-balance KL monotonicity (Proposition 6.4) are used here in their classical finite-state forms and credited explicitly to [18, 19, 20, 21, 22, 11, 23]. Lemma 2.1 and Proposition 3.1 record elementary consequences of the specific generator J . Table 1 lists where each result enters the chain.

The reader who is uninterested in the cost-determining functional equation can read Sections 5–6 as a self-contained treatment of the canonical ensemble on a finite space with an arbitrary fixed cost vector X_ω . The role of Sections 1.2–3 is to show how, in ratio-bookkeeping settings such as detailed balance and chemical equilibrium, the cost vector is selected once one adopts the multiplicative–reciprocal closure and calibration on the underlying ratios.

Recognition temperature and free energy; positive-temperature scope. We refer to the dimensionless quantity $T_R := 1/\beta$, with β the Lagrange multiplier for the mean-cost

constraint, as the recognition temperature, and to $F_R = \langle X \rangle_q - T_R H(q)$ as the recognition free energy. Both names are by analogy with the canonical ensemble: T_R and F_R are the Legendre-conjugate variables to mean cost and entropy, and obey the standard Helmholtz algebra in the cost notation $X_\omega = J(r_\omega)$. T_R has dimensions of cost, not Kelvin, and the identification $T_R = k_B T$ for a concrete material requires a microscopic model that the present finite-state framework does not supply.‡ Sections 5–6 restrict attention to $\beta > 0$ (equivalently E_{mc} below the uniform-cost mean, see Proposition 5.2); the negative-temperature regime $\beta < 0$ is treated alongside the cosine branch of (8) in [13], with a recent H -theorem for bounded-spectrum systems in this regime in [24], and is not pursued here.

Roadmap. Section 2 establishes the quadratic lower bound on J in log-ratio coordinates. Section 3 defines the J -cost divergence D_J and its half-Neyman form. Section 4 recovers Shannon entropy from multinomial counting and proves the soft-shell constrained multinomial theorem (Theorem 4.4), its exact-shell corollary for rational data (Corollary 4.5), and the non-asymptotic Stirling bound of (N2); Section 8.7 records the multi-constraint soft-shell extension (N2-soft). Section 5 produces the finite-state Gibbs law via Lagrange optimization. Section 6 proves the free-energy–Kullback–Leibler identity and the detailed-balance monotonicity of D_{KL} . Section 7 works through a three-state example, including a numerical relaxation trajectory under a detailed-balance Markov dynamics, the Tsallis comparison, and the sample-size power calculation at fixed RCL ground truth (Section 7.6). Section 8 addresses context and limitations, including the connection to stochastic thermodynamics (Section 8.3) and to the maximum-entropy and large-deviation literature; Section 9 closes.

‡ All quantities in Sections 5–6 are dimensionless: r_ω is a positive number, $J(r_\omega)$ is a positive number, and so are X_ω , E_{mc} , β , $T_R = 1/\beta$, Z , F_R , and H . To map this dimensionless framework onto a model in which X_ω represents a physical energy E_ω in joules, fix an energy scale ε_0 (with units of joules) and set $E_\omega := \varepsilon_0 X_\omega$, $T := \varepsilon_0 T_R / k_B$ (so $k_B T = \varepsilon_0 T_R$, a numerical equality between two energies in joules), and $F^{\text{SI}} := \varepsilon_0 F_R$. The Boltzmann factor $e^{-E_\omega / k_B T} = e^{-\beta X_\omega}$ is invariant; all identities of Sections 5–6 carry over unchanged after this rescaling. The bridge constant ε_0 must be supplied by a microscopic model and is not part of the present finite-state framework.

Result	Role in the chain
Theorem 1.1 (classification of J under measurability and local integrability; [13])	Reduces (A1)–(A3) to the d’Alembert functional equation [14, 15] in log-coordinates and selects the hyperbolic branch; supplies the closed cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ used everywhere downstream.
Lemma 2.1 (squared-log lower bound on J)	Hyperbolic-cosine identity $J(e^t) = \cosh t - 1 \geq t^2/2$; feeds the squared-log bound on D_J and the surrogate comparison in the example of Section 7.
Proposition 3.1 (D_J properties; half-Neyman form)	f -divergence with generator $f = J$; supplies a closed Neyman- χ^2 form on the simplex used in the numerical example.
Theorem 4.4, Corollary 4.5, Propositions 4.1, 8.1 (constrained types; exact rational shell; non-asymptotic and multi-constraint soft-shell bounds)	Method-of-types statement on the affine cost shell selected by $X_\omega = J(r_\omega)$, with explicit finite- N, K control via Robbins’ inequalities; soft-shell and multi-constraint extensions for irrational measured costs.
Proposition 5.3 (Gibbs variational principle)	Lagrangian construction in the cost notation $X_\omega = J(r_\omega)$; supplies the canonical Gibbs weights and the dual parameter β .
Proposition 6.1, Corollary 6.2 (free-energy–KL identity)	Algebraic identity $F_R(q) - F_R(p) = T_R D_{\text{KL}}(q p)$ in the recognition variables, with the corresponding minimum-free-energy dual. Remark 6.3 records that the maximum-entropy dual is logically identical to Proposition 5.3.
Proposition 6.4, Corollary 6.5 (monotone D_{KL} under detailed balance)	Closes the chain at the level of detailed-balance relaxation: $t \mapsto D_{\text{KL}}(q(t) p)$ is nonincreasing along reversible Markov dynamics, hence so is $F_R(q(t))$.

Table 1. Where each formal result enters the chain $\{r_\omega\} \rightarrow J \rightarrow X \rightarrow (\text{Gibbs}, F_R, D_{\text{KL}})$. Contributions (N1)–(N3) are summarized in Section 1.3; (N1) and (N3) are expository, (N2) records Robbins’ non-asymptotic bound, and (N2-soft) gives the soft-shell full-sequence theorem in single- and multi-constraint form.

2. The RCL cost J and its quadratic lower bound

For every $x > 0$, the closed-form cost (5) is bounded below by its second-order Taylor expansion in log-coordinates: $J(x) \geq (\log x)^2/2$, with equality iff $x = 1$ (Lemma 2.1). The proof is a one-line consequence of the power series $\cosh t - 1 = \sum_{k \geq 1} t^{2k}/(2k)!$. We record the bound because it is applied pointwise to likelihood ratios in the lower bound on D_J in Section 3, and because it sets up the squared-log surrogate cost vector $X_\omega^{\text{SQ}} = (\log r_\omega)^2/2$ used in the discriminating-observable analysis of Section 7.

Log-coordinate form. In logarithmic coordinates the cost simplifies considerably. Writing any ratio $x > 0$ as $x = e^t$ with $t = \log x$, the RCL cost becomes $J(e^t) = \cosh t - 1$, so multiplicative departures of x from 1 are represented by additive deviations $t = \log x$.

The reciprocal symmetry $J(x) = J(1/x)$ reduces in these coordinates to the evenness $\cosh(-t) = \cosh t$, and the vanishing of J at $x = 1$ to $\cosh 0 - 1 = 0$.

The quadratic lower bound. Since $\cosh t - 1 = \sum_{k \geq 1} t^{2k}/(2k)!$, every Taylor coefficient at $t = 0$ beyond second order is nonnegative, so $\cosh t - 1 \geq t^2/2$ for all real t , with equality only at $t = 0$. Equivalently, $2(\cosh t - 1) = (e^{t/2} - e^{-t/2})^2$ reduces the bound to $|\sinh u| \geq |u|$ for $u = t/2$. In the variable x this is the following statement; Section 3 applies it pointwise to likelihood ratios when bounding D_J (Proposition 3.1).

Lemma 2.1 (Quadratic lower bound on J). *For all $x > 0$,*

$$J(x) \geq \frac{(\log x)^2}{2}, \quad (9)$$

with equality if and only if $x = 1$.

Proof. Set $t = \log x$. Then $J(x) = \cosh t - 1 = \sum_{k \geq 1} t^{2k}/(2k)! \geq t^2/2 = (\log x)^2/2$, with equality if and only if $t = 0$. \square

The bound (9) is used twice downstream: pointwise on likelihood ratios in Proposition 3.1 (iii) to control D_J from below by a p -weighted squared logarithm, and at the level of single ratios in Section 7 to compare the RCL cost vector with the squared-log surrogate $X_\omega^{\text{SQ}} = (\log r_\omega)^2/2$ on the same three-state example (Table 5).

Figure 1 compares the two functions. The RCL cost J and the squared-log surrogate $(\log x)^2/2$ share a common second-order expansion at $x = 1$ (so the bound is sharp there: replacing $1/2$ by any $c > 1/2$ would violate (9) in a neighborhood of $x = 1$). Away from $x = 1$, the gap is the strictly positive tail $\sum_{k \geq 2} t^{2k}/(2k)!$ of the hyperbolic-cosine series; in particular, no constant $c > 1/2$ satisfies $J(x) \geq c(\log x)^2$ for all $x > 0$, because $J(e^t) = \cosh t - 1$ grows like $\frac{1}{2}e^{|t|}$ in the log-coordinate $t = \log x$ (equivalently, $J(x)$ grows linearly in $\max(x, 1/x)$ as a function of the ratio variable x), while $(\log x)^2$ grows only quadratically in $|\log x|$. Thus $J(x)/(\log x)^2 \rightarrow \infty$ as $x \rightarrow \infty$, while $J(x)/x \rightarrow 1/2$: J grows faster than any power of $|\log x|$, but only linearly in the ratio variable x itself.

Interpretation: squared-log fluctuations. Lemma 2.1 is a pointwise comparison between two functions of a single ratio; it does not identify J with Shannon entropy, Fisher information, or a relative-entropy functional. Shannon entropy enters only later, through the multinomial counting of Section 4; it lives on probability vectors, not on ratios, and is supplied by combinatorics rather than by the functional equation (2).

Lemma 2.1 shows that the RCL-selected cost controls the squared logarithmic deviation of each ratio from the reference state. Two regimes are easy to describe. Near equilibrium ($x \approx 1$), $J(x) \approx (x-1)^2/2 \approx (\log x)^2/2$, so any near-equilibrium expansion of a J -based quantity reproduces, at second order, the corresponding squared-log expression. Far from equilibrium ($x \gg 1$ or $x \ll 1$), the strict convexity of $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ penalizes large multiplicative deviations more sharply than the quadratic surrogate, since $J(e^t)$ grows like $\frac{1}{2}e^{|t|}$ in the log-coordinate $t = \log x$ (equivalently, $J(x)$ grows

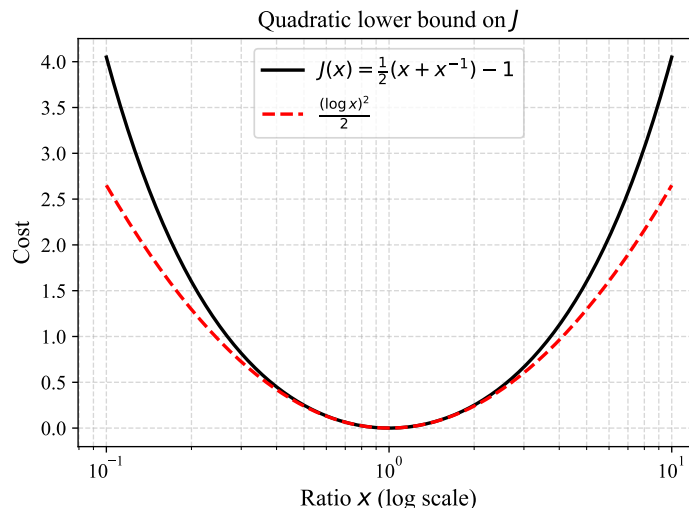


Figure 1. Quadratic lower bound on J . The RCL cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ (solid curve) and its quadratic surrogate $(\log x)^2/2$ (dashed curve). The two curves share a common second-order expansion at $x = 1$; the gap grows for $x \ll 1$ or $x \gg 1$, where $J(e^t)$ grows like $\frac{1}{2}e^{|t|}$ in the log-coordinate $t = \log x$ (equivalently, $J(x)$ grows linearly in $\max(x, 1/x)$), while $(\log x)^2/2$ grows only quadratically in $|t|$.

linearly in $\max(x, 1/x)$), whereas $(\log x)^2/2$ grows only quadratically in $|t|$. Thus J sits above a squared-log fluctuation measure in ratio coordinates, while H remains the combinatorial rate functional for empirical compositions, and any far-from-equilibrium dynamics built on J would penalize large multiplicative deviations more sharply than a quadratic log-ratio surrogate.

3. The J -cost divergence

Throughout this section, Ω denotes a finite set and p, q are strictly positive probability vectors on Ω , i.e. $p, q : \Omega \rightarrow (0, 1]$ with $\sum_{\omega \in \Omega} p(\omega) = \sum_{\omega \in \Omega} q(\omega) = 1$; in this section p plays the role of a generic reference distribution and q that of a generic competitor, with no Gibbs interpretation. From Section 5 onwards, p is reserved for the Gibbs reference (20) (parametric family $p^{(\beta)}$ when the dependence on β is exhibited, and $p^* = p^{(\beta^*)}$ when the constraint-matching value β^* from Proposition 5.2 is fixed), while q continues to denote a generic competitor distribution. Strict positivity avoids boundary issues peripheral to the arguments below; the divergence extends to the boundary by lower-semicontinuity, but that extension is not used here. The Kullback–Leibler (relative-entropy) divergence [25]

$$D_{\text{KL}}(q||p) := \sum_{\omega \in \Omega} q(\omega) \log \frac{q(\omega)}{p(\omega)}$$

will be used at one point in Remark 3.2 for comparison with D_J ; its main role is in Section 6, where it is reintroduced with the same convention $0 \log 0 := 0$.

The pointwise comparison in Lemma 2.1 concerns two functions of a single positive ratio. A probabilistic version requires aggregating across outcomes: at each outcome ω

we form the likelihood ratio $q(\omega)/p(\omega)$, evaluate the RCL cost $J(\cdot)$ at this ratio, and average with respect to the reference distribution p . The result is a divergence in the standard probabilistic sense—a nonnegative functional of the pair (q, p) that vanishes precisely at coincidence—and it is the natural ratio-level companion of the cost fixed in Section 2.

3.1. Definition and half-Neyman χ^2 form

Define the J -cost divergence by

$$D_J(q\|p) = \sum_{\omega \in \Omega} p(\omega) J\left(\frac{q(\omega)}{p(\omega)}\right). \quad (10)$$

This also places J near familiar ratio-based contrasts, since the symmetrized Itakura–Saito divergence on a single positive ratio is proportional to $r + r^{-1} - 2$. This is the standard recipe for converting an even, convex, single-variable cost on positive ratios into a divergence on a finite simplex; in the language of information theory, D_J is the Csiszár f -divergence $D_f(q\|p) = \sum_{\omega} p_{\omega} f(q_{\omega}/p_{\omega})$ with $f = J$ [26, 21, 27]. Here J is not an arbitrary convex f but is fixed by the functional equation (2) together with the unit log-curvature calibration, so the generator in (10) is not a separate modeling choice.

Since $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ is a rational function of x , the sum (10) expands explicitly. Substituting $J(q_{\omega}/p_{\omega}) = \frac{1}{2}(q_{\omega}/p_{\omega} + p_{\omega}/q_{\omega}) - 1 = (q_{\omega} - p_{\omega})^2/(2p_{\omega}q_{\omega})$ into (10) and multiplying by p_{ω} produces the equivalent half-Neyman representation

$$D_J(q\|p) = \frac{1}{2} \sum_{\omega \in \Omega} \frac{(p(\omega) - q(\omega))^2}{q(\omega)}, \quad (11)$$

identifying D_J with one-half of Neyman’s χ^2 divergence $\sum_{\omega} (p_{\omega} - q_{\omega})^2/q_{\omega}$ with arguments oriented as in (10). Here q appears in the denominator of each summand; this is Neyman’s orientation, as opposed to Pearson’s χ^2 , which places p in the denominator [26]. The representation makes it manifest that $D_J \geq 0$ with equality only at $p \equiv q$; furthermore $J''(x) = x^{-3} > 0$ on $\mathbb{R}_{>0}$, so J is strictly convex and the abstract theory of f -divergences applies without modification.

3.2. Positivity and squared-log domination

Lemma 2.1 feeds directly into a lower bound on D_J . Applying (9) pointwise to the likelihood ratio q_{ω}/p_{ω} , multiplying by p_{ω} , and summing over ω yields control of D_J by a squared logarithmic deviation. The following proposition combines the closed-form representation (11) with this bound; both are used in the example of Section 7.

Proposition 3.1 (Basic properties of D_J). *For all positive probability distributions p, q on Ω :*

- (i) $D_J(q\|p) \geq 0$, with equality if and only if $q = p$;

- (ii) $D_J(q||p)$ is given by (11);
 (iii) one has the squared-log lower bound

$$D_J(q||p) \geq \frac{1}{2} \sum_{\omega \in \Omega} p(\omega) (\log(q(\omega)/p(\omega)))^2. \quad (12)$$

Proof. For (i), each summand in D_J is nonnegative because $J \geq 0$, and $J(x) = 0$ if and only if $x = 1$, so the sum vanishes if and only if $q_\omega/p_\omega = 1$ for all ω . Part (ii) is the algebraic expansion already given in (11). For (iii), apply Lemma 2.1 pointwise to $x = q_\omega/p_\omega$, multiply by $p(\omega)$, and sum over $\omega \in \Omega$. \square

The closed form (11) is convenient for numerical work—it is what is used in the three-state example of Section 7—but D_J is not a substitute for the Kullback–Leibler divergence. It does not generate the free-energy identity that organizes Section 6; rather, it is the f -divergence naturally associated with the ratio-level cost J , and Proposition 3.1 (iii) is the divergence-level reflection of Lemma 2.1. The KL geometry that drives the variational principle re-enters at the level of logarithms of probabilities, not at the level of ratios, and Section 6 keeps D_J and D_{KL} formally separate for this reason.

The right-hand side of (12) is itself a squared-log quantity. To leading quadratic order near $q = p$, it is one half of the usual Pearson chi-square quadratic form, $\frac{1}{2} \sum_{\omega} (q_\omega - p_\omega)^2/p_\omega$, but the two expressions are not identical away from coincidence. Proposition 3.1 (iii) provides one-sided pointwise control, not a two-sided comparison with D_{KL} .

3.3. Relation to KL geometry

Both D_J and D_{KL} are f -divergences on the simplex (with $f = J$ and $f(x) = x \log x - x + 1$ respectively), both vanish iff $q = p$, and both are jointly convex in (q, p) . They are nevertheless not interchangeable as ordering rules.

Remark 3.2 (D_J is not interchangeable with D_{KL}). *Distinct f -divergences need not induce the same ordering among alternatives to a fixed reference. For $p = (0.3, 0.7)$, $q^{(1)} = (0.1, 0.9)$, $q^{(2)} = (0.6, 0.4)$ on $|\Omega| = 2$, one finds $D_{\text{KL}}(q^{(1)}||p) \approx 0.116 < D_{\text{KL}}(q^{(2)}||p) \approx 0.192$ but $D_J(q^{(1)}||p) \approx 0.222 > D_J(q^{(2)}||p) \approx 0.188$. Any criterion—gradient flow, model selection, rate-function minimization—that uses one divergence as a contrast may therefore prefer a different alternative under the other.*

The variational principle for the Gibbs distribution (Section 5) and the free-energy identity (Section 6) both pass through D_{KL} , since the log-affine form of the Gibbs measure and Shannon entropy on probability vectors couple to log-probabilities, not to ratios; D_{KL} is the f -divergence aligned with that structure on the simplex. From Section 4 onward, D_{KL} is the operative divergence; D_J serves only as the ratio-level companion of the cost J , where it appears in Proposition 3.1 and in the squared-log comparison of Table 2.

To summarize the role of this section: (11) gives a closed-form numerical expression for D_J that is used directly in Section 7.3, while (12) feeds the squared-log surrogate cost vector X_ω^{SQ} used in the discriminating-observable analysis (Table 5). The KL divergence D_{KL} takes over from Section 4 onward and is the divergence that organizes the free-energy identity in Section 6.

4. Finite-state sample model and entropy from counting

Notation used in Sections 4–6. For convenience we collect the recurring symbols. $\Omega = \{\omega_1, \dots, \omega_K\}$ is a finite outcome space of size $K \geq 2$. For each outcome ω , the positive ratio $r_\omega > 0$ (the input ratio for the modeling problem) determines the cost $X_\omega := J(r_\omega) \in \mathbb{R}$ via the closed form (5). A probability distribution on Ω is denoted p when it is the Gibbs reference measure $p_\omega^{(\beta)} \propto e^{-\beta X_\omega}$ and q when it is a generic alternative. The constraint value $E_{\text{mc}} \in \mathbb{R}$ is the prescribed mean of the cost vector X (target for $\langle X \rangle_p$ under the variational problem); it is not a physical energy, which appears as E_ω only in the SI bridge in the footnote to the recognition-temperature discussion in Section 1.3. The Lagrange multiplier $\beta \in \mathbb{R}$ is conjugate to that mean; $T_{\text{R}} = 1/\beta$ for $\beta > 0$; $Z(\beta) = \sum_\omega e^{-\beta X_\omega}$ is the partition function. The empirical type $\hat{p}_k = n_k/N$ is reserved for finite-sample frequency vectors in this section. H denotes Shannon entropy on Δ_K ; D_{KL} and D_J are as in Section 3; $F_{\text{R}} = \langle X \rangle - T_{\text{R}}H$ is the recognition free energy. Throughout this section and Sections 5–6 we assume X is non-constant on Ω and $\min_\omega X_\omega < E_{\text{mc}} < \max_\omega X_\omega$; without this the Gibbs construction is degenerate.

While J is selected by the composition law and calibration, the Shannon entropy functional used in the next section is not an additional axiom. It is recovered by counting microstates: for an array of N independent outcomes drawn from K types, enumerate the integer compositions $\mathbf{n} = (n_1, \dots, n_K)$ compatible with a prescribed mean of an intensive observable, take the logarithm of the resulting multinomial coefficient, and pass to the large- N limit. The leading term is $NH(\hat{p})$ with $\hat{p}_k = n_k/N$ the empirical frequency vector. The constrained multinomial entropy theorem of this section is the finite-state, method-of-types instance of Sanov’s theorem, in the sense of Cover and Thomas [20]. Density hypotheses on admissible system sizes are required because on a discrete sample space the affine constraint slice may not be reached by integer compositions for every N ; Section 4.3 states the lifting and density properties needed.

4.1. Sample of observations and empirical types

Consider a sample of N independent observations, each reporting a ratio $r_i \in \{r^{(1)}, \dots, r^{(K)}\}$ drawn from K distinct values. The ordered sequence (r_1, \dots, r_N) is a sample, and the unordered tally $\mathbf{n} = (n_1, \dots, n_K)$, with $n_k \geq 0$ counting how many of the N observations report value $r^{(k)}$ and with $\sum_k n_k = N$, is its empirical type. Two samples with the same type differ only by the order of the observations; we count both.

The number of microstates compatible with a given macrostate \mathbf{n} is the multinomial

coefficient

$$W(\mathbf{n}) = \frac{N!}{n_1! n_2! \cdots n_K!}, \quad (13)$$

counting the arrangements of N distinguishable events into K classes of sizes n_1, \dots, n_K . The macrostate supports the empirical distribution $\hat{p}_k = n_k/N \in [0, 1]$ on $\{1, \dots, K\}$.

4.2. Multinomial weights and Shannon entropy

Applying Stirling's approximation $\log(m!) = m \log m - m + O(\log m)$ to each factorial in (13) and rewriting in terms of \hat{p} gives an asymptotic identity in which the implicit constant can be taken uniform over macrostates with $n_k \geq 1$ for all k (boundary cases where some n_k vanishes contribute only lower-order corrections under the convention $0 \log 0 := 0$):

$$\frac{1}{N} \log W(\mathbf{n}) = - \sum_{k=1}^K \hat{p}_k \log \hat{p}_k + O\left(\frac{K \log N}{N}\right) =: H_{\text{type}}(\hat{p}) + O\left(\frac{K \log N}{N}\right). \quad (14)$$

H_{type} denotes Shannon entropy at the empirical frequency vector $\hat{p} \in \Delta_K$. The same functional on a generic finite outcome space Ω is denoted H in Sections 5–6; the subscript “type” is retained only here, where \hat{p} is an integer-rationally-quantized empirical type. The identity (14) is the dimensionless Boltzmann principle $S = k_B \log W$ [1, 18, 20] on the per-event scale, with k_B absorbed into the choice of nat units; it is the standard finite-state derivation of Shannon entropy as the rate function of a multinomial array, restricted to the affine cost shell selected by the constraint of Section 4.3.

The asymptotic statement (14) hides a non-asymptotic Stirling bound that is needed for assessing the regime where K grows with N . We record the explicit bound below; the proof reduces to Robbins' two-sided Stirling inequalities [17] applied to each factorial.

Proposition 4.1 (Non-asymptotic Stirling bound on $\log W$). *For all $N \geq 1$, all $K \geq 2$, and all $\mathbf{n} \in \mathbb{Z}_{\geq 0}^K$ with $\sum_k n_k = N$ and $K^\dagger := \#\{k : n_k \geq 1\}$ the number of nonempty types, the multinomial weight $W(\mathbf{n}) = N! / \prod_k n_k!$ satisfies*

$$\left| \frac{1}{N} \log W(\mathbf{n}) - H(\hat{p}) \right| \leq \frac{K^\dagger - 1}{2N} \log N + \frac{R^*(K^\dagger)}{N}, \quad (15)$$

where

$$R^*(K^\dagger) := \frac{K^\dagger - 1}{2} \log(2\pi) + \frac{1}{12} + \frac{K^\dagger}{12}.$$

In particular, when $K^\dagger \leq K$ is fixed, the right-hand side of (15) is $(K - 1) \log N / (2N) + O(K/N)$, recovering (14) with explicit constants. When K^\dagger grows with N , the bound degrades: (15) is non-trivial only when $K \log N / N \rightarrow 0$, equivalently $N / (K \log N) \rightarrow \infty$.

Proof. Let $S = \{k : n_k \geq 1\}$, so $|S| = K^\dagger$. Robbins' inequalities [17] give, for every $m \geq 1$,

$$\log m! = m \log m - m + \frac{1}{2} \log(2\pi m) + \varepsilon_m, \quad 0 < \varepsilon_m < \frac{1}{12m}. \quad (16)$$

Applying (16) to $N!$ and to $n_k!$ for $k \in S$, and using $\log 0! = 0$ for empty types, gives

$$\log W(\mathbf{n}) = N \log N - \sum_{k \in S} n_k \log n_k + \frac{1}{2} \left(\log(2\pi N) - \sum_{k \in S} \log(2\pi n_k) \right) + \varepsilon_N - \sum_{k \in S} \varepsilon_{n_k}.$$

Since

$$H(\hat{p}) = \log N - \frac{1}{N} \sum_{k \in S} n_k \log n_k,$$

we obtain

$$\frac{1}{N} \log W(\mathbf{n}) - H(\hat{p}) = \frac{1}{2N} \left(\log(2\pi N) - \sum_{k \in S} \log(2\pi n_k) \right) + \frac{1}{N} \left(\varepsilon_N - \sum_{k \in S} \varepsilon_{n_k} \right).$$

The logarithmic term is bounded in absolute value by

$$\frac{K^\dagger - 1}{2N} \log N + \frac{K^\dagger - 1}{2N} \log(2\pi),$$

because (i) $1 \leq n_k \leq N$ for $k \in S$ implies $0 \leq \sum_{k \in S} \log n_k \leq K^\dagger \log N$, so $|\log N - \sum_{k \in S} \log n_k| \leq (K^\dagger - 1) \log N$, and (ii) the constant pieces collect into $|(K^\dagger - 1) \log(2\pi)| = (K^\dagger - 1) \log(2\pi)$. The Robbins remainders satisfy

$$\left| \varepsilon_N - \sum_{k \in S} \varepsilon_{n_k} \right| \leq \frac{1}{12} + \frac{K^\dagger}{12}.$$

Combining these estimates yields the stated bound. \square

For the three-state example of Section 7 ($K = 3$) at $N = 100$, taking $K^\dagger = 3$ gives the leading term $(K^\dagger - 1) \log N / (2N) = \log(100) / 100 \approx 0.046$ and the Robbins remainder $R^*(3) / N = (\log(2\pi) + 1/12 + 1/4) / 100 \approx 0.022$, so the right-hand side of (15) is at most ≈ 0.068 nats per observation. This is a uniform absolute error bound on the Stirling approximation to $N^{-1} \log W(\mathbf{n})$, not a relative error bound on the Gibbs weights or on the rate function in any state-dependent sense; thus the rate-function approximation $\frac{1}{N} \log W(\mathbf{n}) \approx H(\hat{p})$ is controlled to within this additive margin at $N = 100$. The bound is informative in the fixed- K regime $N \gg K \log N$ and degrades when $K \log N$ approaches N (Remark 4.2). As a working rule of thumb in the fixed- K regime, $N \gtrsim 30K \log K$ keeps the right-hand side of (15) below $\log K / 10$.

To leading order in N , the number of microstates supporting a given empirical composition \hat{p} thus grows exponentially at rate $H(\hat{p})$. Without constraints, every composition is admissible and the maximizer of W is the maximizer of H , namely the uniform distribution. Imposing a linear constraint on the empirical mean of an intensive quantity restricts admissible macrostates to a strict subset, and the maximum of H over this subset selects an empirical distribution that is generally non-uniform. The constrained multinomial entropy theorem below is the finite-state formalization of this picture and the bridge to the Gibbs variational principle of Section 5.

Remark 4.2 (Regime of validity: fixed K , $N \rightarrow \infty$). *The asymptotic results of this section, including Theorem 4.4 and Proposition 4.1, are stated in the limit $N \rightarrow \infty$ at fixed finite outcome space Ω (equivalently fixed $K = |\Omega|$). This is the inferential regime in which the number of independent observations grows while the underlying system, including its state count, remains fixed; it differs from the classical thermodynamic limit, in which the system size and hence the state count grow, and from large-deviation regimes where $K = K(N) \rightarrow \infty$ jointly with N . The bound (15) is non-vacuous precisely when $N \gg K \log N$, equivalently $N \gtrsim K \log K$ for the validity of (14) as a leading-order asymptotic, and degrades to vacuity once $K \log N$ approaches N . As a consequence, for models in which K is exponential in a system size L (e.g., $K = 2^L$ for an L -spin system) the regime $N \gg K \log N$ requires $N \gtrsim L 2^L$, far beyond what a realistic sample can support; Theorem 4.4 therefore applies to coarse-grained or effective-state-count models with K fixed (or with the active type count K^\dagger controlled), but not as a tight bound for spatially extended lattice models.*

4.3. Constrained types and admissible system sizes

The unconstrained maximizer of H is uniform; the next two subsections specify the constrained problem (this subsection) and prove that constrained multinomial weights are again maximized by an entropy maximizer, now on a restricted slice of the simplex (Section 4.4). Two ingredients are needed: a definition of the set of integer macrostates compatible with the constraint at fixed system size N , and a verification that this set is nonempty often enough to support a meaningful $N \rightarrow \infty$ limit. The first is straightforward; the second is a Diophantine subtlety addressed by Definition 4.3 below.

Heuristically: the continuous constraint set $\Gamma \subset \Delta_K$ is an affine slice, and rational types $p \in \mathbb{Q}^K \cap \Gamma$ are dense in it; but “ N events with type p ” requires the integers Np_k to actually be integers. The basic arithmetic observation recorded after (17) below gives the lift, and Definition 4.3 specifies the subsequence of N ’s along which lifts of arbitrarily entropy-near-maximal rational types are simultaneously available.

The relevant constraint is the standard one of equilibrium statistical mechanics: a fixed average value of an intensive scalar cost. With cost values $g_k := J(r^{(k)})$, the empirical mean cost at macrostate \mathbf{n} is $\sum_k \hat{p}_k g_k = N^{-1} \sum_k n_k g_k$, and fixing this mean to a target value E_{mc} defines an affine slice of the simplex.

Concretely, fix $g = (g_1, \dots, g_K) \in \mathbb{R}^K$ and $E_{\text{mc}} \in \mathbb{R}$. For each $N \in \mathbb{N}$, define the set of nonnegative integer compositions on the affine cost shell by

$$\mathcal{F}_N(g, E_{\text{mc}}) := \left\{ \mathbf{n} \in \mathbb{Z}_{\geq 0}^K : \sum_{k=1}^K n_k = N, \sum_{k=1}^K n_k g_k = N E_{\text{mc}} \right\}. \quad (17)$$

This is the discrete object that the constrained multinomial theorem of Section 4.4 will count. The pair (g, E_{mc}) is macro-feasible if $\mathcal{F}_N(g, E_{\text{mc}}) \neq \emptyset$ for at least one positive integer N ; equivalently, setting $p_k^\dagger = n_k/N$ for any feasible \mathbf{n} , there exists a rational type $p^\dagger \in \Delta_K \cap \mathbb{Q}^K$ with $\sum_k p_k^\dagger g_k = E_{\text{mc}}$. Macro-feasibility is therefore a Diophantine condition on (g, E_{mc}) , not a geometric one on the constraint slice.

Although the continuous constraint slice $\Gamma = \{p \in \Delta_K : \sum_k p_k g_k = E_{\text{mc}}\}$ may be a nonempty subset of the simplex, the integer equation $\sum_k n_k g_k = N E_{\text{mc}}$ need not be solvable for every N , so $\mathcal{F}_N(g, E_{\text{mc}})$ can be empty. A naive limit $N \rightarrow \infty$ may wander through values of N at which there is nothing to count, so the exact-shell asymptotic statement in Corollary 4.5 below is phrased along subsequences for which integer compositions of the right type both exist and accumulate near the entropy maximizer. The relevant lifting and density properties are stated in Definition 4.3 below; this is the only Diophantine input to the corollary. As a basic arithmetic observation: if $q \in \Delta_K \cap \mathbb{Q}^K$ has common denominator $Q \in \mathbb{N}$ (so $Qq_k \in \mathbb{Z}_{\geq 0}$ for every k) and $\sum_k q_k g_k = E_{\text{mc}}$ with $g \in \mathbb{Q}^K, E_{\text{mc}} \in \mathbb{Q}$, then for any N divisible by Q the composition $n_k := Nq_k$ is an element of $\mathcal{F}_N(g, E_{\text{mc}})$, because $\sum_k n_k = N \sum_k q_k = N$ and $\sum_k n_k g_k = N \sum_k q_k g_k = N E_{\text{mc}}$.

Definition 4.3 (Denominator-exhaustive sequence). *With $h_\star := \max_{p \in \Gamma} H(p)$ and $\mathcal{I} := \{N \in \mathbb{N} : \mathcal{F}_N \neq \emptyset\}$, a sequence $(N_j)_{j \geq 1} \subset \mathcal{I}$ is denominator-exhaustive if there exists a countable family $\{q^{(m)}\}_{m \geq 1} \subset \Gamma \cap \mathbb{Q}^K$ with $H(q^{(m)}) \rightarrow h_\star$ as $m \rightarrow \infty$ and with common denominators $Q^{(m)}$ (so $Q^{(m)} q_k^{(m)} \in \mathbb{Z}_{\geq 0}$ for every k) such that, for every fixed m , $Q^{(m)}$ divides N_j for all sufficiently large j (i.e. there exists $J(m) \in \mathbb{N}$ with $Q^{(m)} \mid N_j$ for every $j \geq J(m)$).*

Eventual divisibility is the strengthening relative to the weaker “divides infinitely many N_j ” formulation: it ensures that the lifted compositions $\mathbf{n}^{(j,m)} = N_j q^{(m)} \in \mathcal{F}_{N_j}$ are available for every $j \geq J(m)$, so the lower bound below can be obtained along the full tail of the sequence rather than only along a sparser sub-subsequence. Concrete admissible patterns include $N_j = j!$ (eventually divisible by every fixed denominator) and arithmetic progressions $N_j = jQ_0$ chosen so that Q_0 is a common multiple of the denominators in an approximating family $\{q^{(m)}\}$. For rational (g, E_{mc}) , periodicity of \mathcal{I} can also be read from the Smith normal form of the $2 \times K$ constraint matrix [28]; this observation is not needed elsewhere.

4.4. Constrained multinomial theorem

Among macrostates compatible with a prescribed average of the costs $g_k = J(r^{(k)})$, the multinomial weight is asymptotically maximized by empirical compositions that maximize Shannon entropy on the corresponding affine slice. The principal asymptotic statement is Theorem 4.4 below, which allows arbitrary real (g, E_{mc}) by relaxing the hard equality $\sum_k n_k g_k = N E_{\text{mc}}$ to a shrinking tolerance window (the *soft shell*). The exact integer equality case for rational data is recovered as Corollary 4.5. The proofs use Stirling’s formula uniformly in types together with continuity of H on the simplex [20, 21, 29, 22, 11], in the form needed for Section 5.

For real-valued (g, E_{mc}) , define the soft-shell feasible set

$$\mathcal{F}_N^{\delta_N}(g, E_{\text{mc}}) := \left\{ \mathbf{n} \in \mathbb{Z}_{\geq 0}^K : \sum_k n_k = N, \left| N^{-1} \sum_k n_k g_k - E_{\text{mc}} \right| \leq \delta_N \right\}, \quad (18)$$

with tolerance $\delta_N > 0$.

Theorem 4.4 (Constrained multinomial weights; soft shell). *Fix $K \geq 2$, $g = (g_1, \dots, g_K) \in \mathbb{R}^K$, and $E_{\text{mc}} \in \mathbb{R}$. Assume the affine slice*

$$\Gamma := \left\{ p \in \Delta_K : \sum_{k=1}^K p_k g_k = E_{\text{mc}} \right\}$$

is nonempty, and set $h_\star := \max_{p \in \Gamma} H(p)$. Let $\delta_N > 0$ satisfy $\delta_N \rightarrow 0$ and $N\delta_N \rightarrow \infty$ —the minimal tolerance scaling under which the argument below yields the limit (19) along the full sequence $N \rightarrow \infty$ (rather than merely along favorable subsequences). Then $\mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})$ is nonempty for all sufficiently large N , and

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \max_{\mathbf{n} \in \mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})} W(\mathbf{n}) = h_\star. \quad (19)$$

Proof. For $\delta \geq 0$, write

$$\Gamma_\delta := \left\{ p \in \Delta_K : |\langle g \rangle_p - E_{\text{mc}}| \leq \delta \right\}, \quad \langle g \rangle_p := \sum_k p_k g_k.$$

The sets Γ_δ are compact and decrease to Γ as $\delta \downarrow 0$. Since H is continuous on the compact simplex,

$$\sup_{p \in \Gamma_\delta} H(p) \longrightarrow h_\star \quad (\delta \downarrow 0).$$

Indeed, if this failed, there would be $\eta > 0$, $\delta_m \downarrow 0$, and $p^{(m)} \in \Gamma_{\delta_m}$ with $H(p^{(m)}) \geq h_\star + \eta$; a convergent subsequence would have a limit $p^\infty \in \Gamma$ by compactness and continuity of $\langle g \rangle_p$, and continuity of H would give $H(p^\infty) \geq h_\star + \eta$, contradicting the definition of h_\star .

For the upper bound, any $\mathbf{n} \in \mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})$ has empirical type $\hat{p} = \mathbf{n}/N \in \Gamma_{\delta_N}$. Proposition 4.1 gives, uniformly in \mathbf{n} at fixed K ,

$$\frac{1}{N} \log W(\mathbf{n}) \leq H(\hat{p}) + o(1) \leq \sup_{p \in \Gamma_{\delta_N}} H(p) + o(1) = h_\star + o(1),$$

so the limsup in (19) is at most h_\star .

For the lower bound, choose $p^\star \in \Gamma$ with $H(p^\star) = h_\star$; such a maximizer exists by compactness of Γ . For each N , choose an integer vector $\mathbf{n}^{(N)}$ with $\sum_k n_k^{(N)} = N$ and

$$\left\| \frac{\mathbf{n}^{(N)}}{N} - p^\star \right\|_1 \leq \frac{K}{N},$$

obtained, for example, by taking floors of Np_k^\star and distributing the remaining mass among the largest fractional parts. Then

$$\left| \langle g \rangle_{\mathbf{n}^{(N)}/N} - E_{\text{mc}} \right| = \left| \langle g \rangle_{\mathbf{n}^{(N)}/N} - \langle g \rangle_{p^\star} \right| \leq \|g\|_\infty \left\| \frac{\mathbf{n}^{(N)}}{N} - p^\star \right\|_1 \leq \frac{K\|g\|_\infty}{N}.$$

Since $N\delta_N \rightarrow \infty$, the last quantity is at most δ_N for all sufficiently large N . Thus $\mathbf{n}^{(N)} \in \mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})$ eventually. By continuity of H and Proposition 4.1,

$$\frac{1}{N} \log W(\mathbf{n}^{(N)}) = H(\mathbf{n}^{(N)}/N) + o(1) \longrightarrow H(p^*) = h_\star.$$

This gives the matching liminf and proves (19). \square

Corollary 4.5 (Exact affine shell under rational costs). *Fix $K \geq 2$ and ratio values $r^{(1)}, \dots, r^{(K)}$ with $g_k := J(r^{(k)})$ not all equal. Write $\mathcal{F}_N := \mathcal{F}_N(g, E_{\text{mc}})$ as in (17). Assume (g, E_{mc}) is macro-feasible, that $\min_k g_k < E_{\text{mc}} < \max_k g_k$, and that $g_1, \dots, g_K, E_{\text{mc}} \in \mathbb{Q}$, so that $\Gamma := \{p \in \Delta_K : \sum_k p_k g_k = E_{\text{mc}}\}$ is a rational affine slice and $\Gamma \cap \mathbb{Q}^K$ is dense in Γ . Set*

$$h_\star := \max_{p \in \Gamma} H(p),$$

and let $\mathcal{I} := \{N \in \mathbb{N} : \mathcal{F}_N \neq \emptyset\}$. Let $(N_j)_{j \geq 1} \subset \mathcal{I}$ be a denominator-exhaustive sequence (Definition 4.3). For each N_j , let $\mathbf{n}^\star \in \mathcal{F}_{N_j}$ maximize $W(\mathbf{n})$ on \mathcal{F}_{N_j} , and set $\hat{p}_k^{(N_j)} = n_k^\star / N_j$. Then every subsequential limit \hat{p}^∞ of $(\hat{p}^{(N_j)})$ in Δ_K satisfies $H(\hat{p}^\infty) = h_\star$, and

$$\lim_{j \rightarrow \infty} \frac{1}{N_j} \log \max_{\mathbf{n} \in \mathcal{F}_{N_j}} W(\mathbf{n}) = h_\star.$$

Proof. Fix any sequence $\delta_N \downarrow 0$ with $N\delta_N \rightarrow \infty$. For every N and every $\mathbf{n} \in \mathcal{F}_N$, the exact equality $\sum_k n_k g_k = N E_{\text{mc}}$ implies $\mathbf{n} \in \mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})$, hence

$$\max_{\mathbf{n} \in \mathcal{F}_N} W(\mathbf{n}) \leq \max_{\mathbf{n} \in \mathcal{F}_N^{\delta_N}(g, E_{\text{mc}})} W(\mathbf{n}).$$

Applying Theorem 4.4 along the subsequence (N_j) gives

$$\limsup_{j \rightarrow \infty} \frac{1}{N_j} \log \max_{\mathbf{n} \in \mathcal{F}_{N_j}} W(\mathbf{n}) \leq h_\star.$$

For the lower bound, write $H_{N_j} := \max_{\mathbf{n} \in \mathcal{F}_{N_j}} \frac{1}{N_j} \log W(\mathbf{n})$ and fix $\varepsilon > 0$. Under the rationality hypothesis, Γ is a rational polytope with nonempty relative interior whenever $\min_k g_k < E_{\text{mc}} < \max_k g_k$, so $\Gamma \cap \mathbb{Q}^K$ is dense in Γ and continuity of H on Δ_K guarantees the existence of a rational near-maximizer $q^{(\varepsilon)} \in \Gamma \cap \mathbb{Q}^K$ with $H(q^{(\varepsilon)}) \geq h_\star - \varepsilon$. By Definition 4.3, choosing m large enough that $H(q^{(m)}) \geq h_\star - \varepsilon$, the denominator $Q^{(m)}$ divides N_j for all $j \geq J(m)$, so for every such j the lifted composition $\mathbf{n}^{(j)} := N_j q^{(m)} \in \mathcal{F}_{N_j}$ has $\tilde{p}^{(j)} = \mathbf{n}^{(j)} / N_j \equiv q^{(m)}$ exactly. Continuity of H gives $H(\tilde{p}^{(j)}) = H(q^{(m)})$, and Stirling's formula yields, for all $j \geq J(m)$,

$$\frac{1}{N_j} \log W(\mathbf{n}^{(j)}) = H(q^{(m)}) + o(1) \geq h_\star - \varepsilon + o(1).$$

Since \mathbf{n}^\star maximizes W on \mathcal{F}_{N_j} , $H_{N_j} \geq h_\star - \varepsilon + o(1)$ for all $j \geq J(m)$. Because eventual divisibility holds at every m , letting $\varepsilon_m \downarrow 0$ along the family $\{q^{(m)}\}$ gives $\liminf_{j \rightarrow \infty} H_{N_j} \geq h_\star$ along the full sequence (N_j) . Combined with the upper bound, $\lim_j H_{N_j} = h_\star$. Any subsequential limit of maximizers attains h_\star because H is continuous on Δ_K (with $0 \log 0 := 0$). \square

Remark 4.6. *Theorem 4.4 is the full-sequence soft-shell method-of-types statement for arbitrary real (g, E_{mc}) with nonempty Γ . Corollary 4.5 is the exact-shell specialization: it holds along any sequence $(N_j) \subset \mathcal{I}$ for which there exist $\mathbf{n}^{(j)} \in \mathcal{F}_{N_j}$ with $H(\mathbf{n}^{(j)})/N_j \rightarrow h_*$, and denominator-exhaustive sequences (Definition 4.3) supply the Diophantine lower bound in the rational case. Without such an attainability hypothesis, sparse subsequences exist for which feasible types fail to approach the entropy maximizer, and convergence along the full sequence $N \rightarrow \infty$ with $\mathcal{F}_N \neq \emptyset$ cannot be claimed for the exact shell.*

Theorem 4.4 is the finite- K method-of-types statement that maximizing the multinomial weight at approximately fixed average cost is equivalent to maximizing Shannon entropy on the constraint slice [20, 21, 29, 22]. Corollary 4.5 packages the classical exact rational slice; the eventual-divisibility bookkeeping of Definition 4.3 is needed only there. Quantitative finite- N , finite- K control is given by Proposition 4.1.

In summary, the cost functional J is fixed by the RCL once the ratios $r^{(k)}$ are specified, the entropy functional H is supplied by combinatorial counting under Stirling's formula, and the constrained variational problem of maximizing $H(p)$ over $\Gamma = \{p \in \Delta_K : \langle g \rangle_p = E_{\text{mc}}\}$ is the asymptotic shadow of maximizing $W(\mathbf{n})$ over $\mathcal{F}_N(g, E_{\text{mc}})$. Section 5 converts this equivalence into an explicit Gibbs measure via Lagrange optimization.

4.5. Modeling inputs: ratios, costs, and physical interpretation

The finite-state framework treats the microscopic labels r_ω as inputs: they are supplied by a ratio-cost model rather than derived from a Hamiltonian or any other underlying dynamical postulate. Each outcome ω is assigned a positive dimensionless ratio r_ω relative to a chosen reference scale, and the RCL then fixes the scalar cost as $X_\omega = J(r_\omega)$.

Once the ratios are specified, the energy-like cost in the canonical variational problem is determined; mapping the ratios r_ω to a material Hamiltonian, a transition rate, an affinity, or an experimentally measurable imbalance is model-specific and lies outside the present scope. In a reversible two-state chain, detailed balance gives the input ratio $r_b = W_{ab}/W_{ba} = \pi_b/\pi_a$ [23], so the corresponding cost vector is $(0, J(r_b))$; the three-state example of Section 7 illustrates the construction in a nontrivial setting, including the Gibbs distribution at a prescribed mean cost and its comparison with the chain's stationary law.

5. Variational derivation of the finite-state Gibbs law

The previous sections supply two ingredients: the closed-form cost $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ from Theorem 1.1, and Shannon entropy as the asymptotic rate function for multinomial weights on the affine cost shell (Theorem 4.4). The convex-analytic step assembles H and $X_\omega = J(r_\omega)$ into a Lagrangian, identifies its unique critical point on the cost shell, and identifies that critical point as the canonical Gibbs measure. The argument is the standard convex duality of the canonical ensemble, applied here with X_ω selected

upstream from the input ratios once the RCL assumptions are adopted. The section makes no use of the explicit form of J from (5): it works for any non-constant cost vector X_ω on a finite outcome space, and the role of the RCL is only to fix that vector. A reader interested only in the canonical ensemble on a finite space may skip the dependence on J and read X as a generic real-valued observable.

5.1. Setup: entropy, mean cost, and inverse temperature

Standing assumptions for Sections 5–6: Ω is a finite set with $|\Omega| \geq 2$; positive ratios $r_\omega > 0$ are specified for each $\omega \in \Omega$; the cost is $X_\omega := J(r_\omega)$; and X is non-constant on Ω (i.e., not all X_ω are equal). The constraint value E_{mc} satisfies $\min_\omega X_\omega < E_{\text{mc}} < \max_\omega X_\omega$. The thermodynamic convention $T_{\text{R}} > 0$ (equivalently $\beta > 0$) is in force, so E_{mc} lies below the uniform-distribution mean $\langle X \rangle_{p^{(0)}}$ (Proposition 5.2).

Two functionals on this set play a central role. The first is Shannon entropy,

$$H(q) := - \sum_{\omega \in \Omega} q_\omega \log q_\omega,$$

written with the same symbol H as in Section 4 (there applied to empirical types $\hat{p} \in \Delta_K$). The second is the expected cost,

$$\langle X \rangle_q := \sum_{\omega \in \Omega} q_\omega X_\omega,$$

which plays the role of $\sum_k \hat{p}_k g_k$ from Section 4. The variational problem is to maximize $H(q)$ subject to a fixed value of $\langle X \rangle_q = E_{\text{mc}}$ together with the normalization $\sum_\omega q_\omega = 1$.

The inverse temperature β is the Lagrange multiplier conjugate to the mean-cost constraint. In the canonical (positive-temperature) regime $\beta > 0$, set the recognition temperature $T_{\text{R}} := 1/\beta$ (dimensionless unless an additional physical scale is supplied). The candidate maximizer then takes the familiar Gibbs form

$$p_\omega = \frac{e^{-X_\omega/T_{\text{R}}}}{Z}, \quad Z = \sum_{\omega' \in \Omega} e^{-X_{\omega'}/T_{\text{R}}}, \quad (20)$$

where Z is the partition function. The next two subsections verify that (20) is the unique maximizer and that T_{R} is uniquely determined by the constraint value E_{mc} .

5.2. Existence and uniqueness of the dual parameter

Working with β rather than T_{R} allows a uniform treatment of the positive- and negative-temperature regimes. For each $\beta \in \mathbb{R}$, define the partition function

$$Z(\beta) := \sum_{\omega \in \Omega} e^{-\beta X_\omega}$$

and the parametric Gibbs law $p_\omega^{(\beta)} := e^{-\beta X_\omega} / Z(\beta)$. The family $\beta \mapsto p^{(\beta)}$ is the exponential family canonically associated with the sufficient statistic X , and the cumulant generating function $\log Z(\beta)$ encodes its differential structure.

Lemma 5.1 (Strict convexity of $\log Z$). *Under the standing assumption that X is non-constant on Ω , $\log Z$ is strictly convex on \mathbb{R} , and*

$$\frac{d}{d\beta} \log Z(\beta) = -\langle X \rangle_{p^{(\beta)}}, \quad \frac{d^2}{d\beta^2} \log Z(\beta) = \text{Var}_{p^{(\beta)}}(X) > 0.$$

Consequently $\beta \mapsto \langle X \rangle_{p^{(\beta)}}$ is strictly decreasing, with

$$\lim_{\beta \rightarrow -\infty} \langle X \rangle_{p^{(\beta)}} = \max_{\omega \in \Omega} X_\omega, \quad \lim_{\beta \rightarrow +\infty} \langle X \rangle_{p^{(\beta)}} = \min_{\omega \in \Omega} X_\omega.$$

Proof. Differentiate $\log Z$ in β ; the variance formula is the standard exponential-family identity [18]. Strict positivity of the variance for non-constant X implies strict convexity of $\log Z$ and hence a strictly decreasing derivative $-\langle X \rangle_{p^{(\beta)}}$. The limiting values follow from the concentration of $p^{(\beta)}$ on the argmax (respectively argmin) of X as $\beta \rightarrow -\infty$ (respectively $\beta \rightarrow +\infty$). \square

As in the cumulant-generating-function picture for exponential families [18], $\log Z$ is a strictly convex potential whose Legendre transform encodes the entropy–mean-cost tradeoff. Strict convexity makes the conjugate variable $\langle X \rangle_{p^{(\beta)}}$ a strictly monotone function of β , taking each value in the open interval $(\min X, \max X)$ exactly once. Each admissible constraint value E_{mc} therefore corresponds to a unique Lagrange multiplier β , as stated in the proposition below.

Proposition 5.2 (Existence and uniqueness of the dual parameter β). *Assume X is not constant on Ω . For every E_{mc} strictly between $\min_{\omega} X_\omega$ and $\max_{\omega} X_\omega$, there exists a unique $\beta \in \mathbb{R}$ such that $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$. Moreover, $\beta > 0$ if and only if $E_{\text{mc}} < \langle X \rangle_{p^{(0)}}$ (the mean cost under the uniform distribution), and $\beta < 0$ if and only if the reverse strict inequality holds. In the thermodynamic convention $T_{\text{R}} > 0$ used in Sections 5–6, we restrict attention to $\beta > 0$, equivalently E_{mc} below the uniform average.*

Proof. By Lemma 5.1, $\beta \mapsto \langle X \rangle_{p^{(\beta)}}$ is continuous and strictly decreasing with the stated limits, hence bijects \mathbb{R} onto $(\min X, \max X)$. The sign of β is determined by comparing E_{mc} to $\langle X \rangle_{p^{(0)}}$, which is the uniform average. \square

5.3. Entropy maximization and Gibbs weights

With β uniquely fixed by Proposition 5.2, the variational principle can be stated and proved. The argument is the classical Lagrange-multiplier calculation for maximum entropy at fixed mean cost: the candidate maximizer produced by stationarity coincides with the parametric Gibbs law $p^{(\beta)}$ from Lemma 5.1, and strict concavity of Shannon entropy promotes this critical point to the unique global maximizer.

Proposition 5.3 (Finite-state Gibbs variational principle; classical finite-state form). *Assume X is not constant on Ω . Let E_{mc} lie strictly between $\min_{\omega} X_\omega$ and $\max_{\omega} X_\omega$, and let $\beta \in \mathbb{R}$ be the unique dual parameter from Proposition 5.2 such that $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$. Among all probability distributions on Ω with $\langle X \rangle_q = E_{\text{mc}}$ and strictly positive weights,*

the Gibbs law $p_\omega^{(\beta)} = e^{-\beta X_\omega} / Z(\beta)$ uniquely maximizes the Shannon entropy $H(q)$. In the thermodynamic convention $T_R > 0$ used in (20) and Section 6, we restrict attention to $\beta > 0$, equivalently $E_{\text{mc}} < \langle X \rangle_{p^{(0)}}$; in this regime $p^{(\beta)}$ coincides with (20) for $T_R = 1/\beta$.

Proof. Form the Lagrangian

$$\mathcal{L}(q) = H(q) - \beta \left(\sum_{\omega} q_{\omega} X_{\omega} - E_{\text{mc}} \right) - \lambda \left(\sum_{\omega} q_{\omega} - 1 \right), \quad (21)$$

where λ is the multiplier conjugate to the normalization constraint. Stationarity in each q_{ω} yields the equations $-\log q_{\omega} - 1 - \beta X_{\omega} - \lambda = 0$, i.e. $\log q_{\omega} = -\beta X_{\omega} - (1 + \lambda)$ for every $\omega \in \Omega$. Setting $c := 1 + \lambda$ (a single ω -independent constant absorbing the multiplier and the constant -1 from $\partial H / \partial q_{\omega}$) and exponentiating gives

$$q_{\omega} = e^{-\beta X_{\omega} - c}, \quad c \in \mathbb{R};$$

imposing $\sum_{\omega} q_{\omega} = 1$ then forces $e^{-c} = 1/Z(\beta)$ and hence $q_{\omega} = e^{-\beta X_{\omega}} / Z(\beta) = p_{\omega}^{(\beta)}$. Strict concavity of H on the interior of the simplex then implies that this interior critical point is the unique maximizer among strictly positive distributions q with $\langle X \rangle_q = E_{\text{mc}}$ [20]; the extension to the full closed simplex (where some q_{ω} may vanish) when X is non-constant and $E_{\text{mc}} \in (\min X, \max X)$ follows from the strict concavity of H on Δ_{Ω} together with the fact that H is continuous on the closed simplex (under the convention $0 \log 0 := 0$): an interior critical point of a strictly concave continuous function on a compact convex set is the unique global maximum on that set whenever the constraint slice has nonempty interior [18, §II.4]. Uniqueness of β itself is Proposition 5.2. \square

Proposition 5.3 is the classical Gibbs variational principle on a finite space, transcribed into the cost notation used here. The Gibbs weights are therefore not introduced as an additional assumption: they arise as the unique interior critical point of a strictly concave functional on a linear constraint slice. Lemma 5.1 is the cumulant-generating-function viewpoint on the canonical ensemble [18], with the identification $T_R = 1/\beta$ used throughout in the positive-temperature regime.

The model-specific input is the cost vector $X_{\omega} = J(r_{\omega})$ fixed by axioms (A1)–(A3) from the ratios $\{r_{\omega}\}$. Given X , the Gibbs law at prescribed mean cost and the free-energy–Kullback–Leibler identities of Section 6 are the standard finite-state objects determined by that cost; Section 8 explains how this assignment fits classical maximum-entropy and large-deviation templates.

6. Free energy, KL identity, and relaxation

The Gibbs law of Section 5 also admits a second, equally fundamental characterization: it is the unique minimizer of a thermodynamic free-energy functional, and the gap between the free energy at an arbitrary distribution and at the Gibbs reference is proportional to a Kullback–Leibler divergence. This identity is purely algebraic—a direct expansion of definitions—and holds for any probability distribution q on Ω ,

independently of any variational principle. The connection to relaxation dynamics enters through Proposition 6.4, which provides a master equation for which $D_{\text{KL}}(\cdot\|p)$ decreases monotonically.

The free-energy–Kullback–Leibler identity (Proposition 6.1) and the Gibbs minimization property (Corollary 6.2) are standard finite-state statements of equilibrium statistical mechanics [18, 20, 11]; the master-equation monotonicity of D_{KL} under detailed balance (Proposition 6.4) is also classical [19, 20, 30]. The proofs given below fix conventions in the cost notation $X_\omega = J(r_\omega)$ used in Sections 7 and 8.

6.1. Free energy and relative entropy

The recognition free energy at dimensionless temperature $T_{\text{R}} > 0$, in direct analogy with the Helmholtz functional $F = E - TS$ of equilibrium thermodynamics [18], is

$$F_{\text{R}}(q) = \langle X \rangle_q - T_{\text{R}} H(q) = \sum_{\omega} q_{\omega} X_{\omega} + T_{\text{R}} \sum_{\omega} q_{\omega} \log q_{\omega},$$

the standard combination of mean cost and temperature-weighted entropy, with the cost $X_\omega = J(r_\omega)$ supplied by the RCL. The directed Kullback–Leibler divergence [25] between probability distributions q and p with $p_\omega > 0$ for all ω is

$$D_{\text{KL}}(q\|p) = \sum_{\omega} q_{\omega} \log \frac{q_{\omega}}{p_{\omega}},$$

under the convention $0 \cdot \log 0 = 0$. The orientation $D_{\text{KL}}(q\|p)$, with q in the first argument, is the orientation that arises when $F_{\text{R}}(q)$ is expanded around a Gibbs reference p , and is kept fixed below.

6.2. Free-energy–KL identity

The identity that ties F_{R} and D_{KL} together is purely algebraic. It depends only on the log-affinity of the Gibbs measure in the cost X and requires no appeal to the variational principle of Section 5.

Proposition 6.1 (Free Energy–KL Identity; classical, see e.g. [18, 20, 11]). *Fix $T_{\text{R}} > 0$, let p denote the Gibbs measure (20) with this T_{R} and partition function Z , and let F_{R} be defined with the same T_{R} . Then for any probability distribution q on Ω ,*

$$F_{\text{R}}(q) - F_{\text{R}}(p) = T_{\text{R}} D_{\text{KL}}(q\|p), \quad (22)$$

where $F_{\text{R}}(p) = -T_{\text{R}} \log Z$.

Proof. By definition, $F_{\text{R}}(q) = \sum_{\omega} q_{\omega} X_{\omega} + T_{\text{R}} \sum_{\omega} q_{\omega} \log q_{\omega}$. Since $p_{\omega} = e^{-X_{\omega}/T_{\text{R}}}/Z$ by (20), taking logarithms gives $\log p_{\omega} = -X_{\omega}/T_{\text{R}} - \log Z$, so $-T_{\text{R}} \log p_{\omega} = X_{\omega} + T_{\text{R}} \log Z$. Substituting this identity into the definition of $F_{\text{R}}(p)$ gives

$$F_{\text{R}}(p) = \sum_{\omega} p_{\omega} X_{\omega} + T_{\text{R}} \sum_{\omega} p_{\omega} \log p_{\omega} = \sum_{\omega} p_{\omega} X_{\omega} - \sum_{\omega} p_{\omega} X_{\omega} - T_{\text{R}} \log Z = -T_{\text{R}} \log Z.$$

Substituting,

$$\begin{aligned} T_{\text{R}} D_{\text{KL}}(q||p) &= T_{\text{R}} \sum_{\omega} q_{\omega} (\log q_{\omega} - \log p_{\omega}) \\ &= T_{\text{R}} \sum_{\omega} q_{\omega} \log q_{\omega} + \sum_{\omega} q_{\omega} X_{\omega} + T_{\text{R}} \log Z \\ &= F_{\text{R}}(q) + T_{\text{R}} \log Z = F_{\text{R}}(q) - F_{\text{R}}(p), \end{aligned}$$

using the identity $F_{\text{R}}(p) = -T_{\text{R}} \log Z$ in the last step. \square

The proof is a direct expansion: the cost-linear terms reorganize because $\log p$ is affine in X for the Gibbs reference, and the remainder assembles into D_{KL} . Proposition 6.1 is therefore a structural identity for the exponential family on Ω , requiring no convex optimization. Two immediate consequences are recorded below: free-energy minimization and, as a remark, the equivalent maximum-entropy statement at fixed cost.

6.3. Variational consequences

The first corollary is the free-energy minimization property: the Gibbs measure p minimizes F_{R} , with the gap to any competitor given exactly by $T_{\text{R}} D_{\text{KL}}$. This is the dual of Proposition 5.3 and a one-line consequence of Proposition 6.1.

Corollary 6.2 (Gibbs minimizes free energy). $F_{\text{R}}(p) \leq F_{\text{R}}(q)$ for all probability distributions q , with equality if and only if $q = p$.

Proof. By Proposition 6.1, $F_{\text{R}}(q) - F_{\text{R}}(p) = T_{\text{R}} D_{\text{KL}}(q||p) \geq 0$, since $T_{\text{R}} > 0$ and $D_{\text{KL}} \geq 0$. The KL divergence vanishes if and only if $q = p$. \square

The maximum-entropy statement of Section 5 can also be recovered directly from Corollary 6.2, without invoking Lagrange stationarity: at fixed mean cost the only term in F_{R} that varies is the entropy term, and the sign of that variation is fixed by the free-energy minimization property.

Remark 6.3 (Same maximum-entropy statement via the free-energy–KL identity). Among all q with $\langle X \rangle_q = \langle X \rangle_p$, $F_{\text{R}}(q) - F_{\text{R}}(p) = (\langle X \rangle_q - T_{\text{R}} H(q)) - (\langle X \rangle_p - T_{\text{R}} H(p)) = -T_{\text{R}}(H(q) - H(p))$, and Corollary 6.2 gives $F_{\text{R}}(q) \geq F_{\text{R}}(p)$, hence $H(q) \leq H(p)$ with equality iff $q = p$. This is Proposition 5.3 re-derived through the free-energy–KL identity rather than by direct Lagrange stationarity, and is logically identical to Proposition 5.3; it is recorded here only as a consistency check on the algebra of Section 6.

6.4. Detailed-balance relaxation and monotonicity

The variational principle is static. Section 5 shows that the Gibbs measure is the entropy maximum at fixed mean cost; Proposition 6.1 above shows that it is also the unique global minimum of F_{R} on the simplex; neither statement involves a dynamics. To couple the static identity to a process, we now consider a continuous-time Markov chain on

Ω whose stationary law is p and whose transition rates satisfy detailed balance with respect to p . Along such a flow, the Kullback–Leibler divergence to the stationary distribution is a Lyapunov function, and Proposition 6.1 translates that monotonicity into the corresponding statement for F_R .

Proposition 6.4 (Monotonicity of D_{KL} under detailed balance; classical, cf. [19, 30, 20]). *Let $Q = (Q_{ij})_{i,j \in \Omega}$ satisfy $Q_{ij} \geq 0$ for $i \neq j$, row sums $\sum_j Q_{ij} = 0$, and detailed balance $p_i Q_{ij} = p_j Q_{ji}$ with respect to a fixed Gibbs measure $p > 0$ on Ω . Suppose $t \mapsto q(t)$ is a differentiable curve in the interior of the simplex $\Delta(\Omega)$ that satisfies the master equation*

$$\dot{q}_j(t) = \sum_{i \in \Omega} q_i(t) Q_{ij}, \quad j \in \Omega. \quad (23)$$

Define $x_i(t) := q_i(t)/p_i$. Then

$$\frac{d}{dt} D_{\text{KL}}(q(t) \| p) = -\frac{1}{2} \sum_{i,j \in \Omega} p_i Q_{ij} (x_j(t) - x_i(t)) (\log x_j(t) - \log x_i(t)) \leq 0, \quad (24)$$

so $t \mapsto D_{\text{KL}}(q(t) \| p)$ is nonincreasing. If, in addition, Q is irreducible on Ω (equivalently, the off-diagonal graph $\{(i, j) : i \neq j, Q_{ij} > 0\}$ is connected as an undirected graph), then by the Perron–Frobenius theorem applied to the rate matrix [31, Theorem 3.5.5], p is the unique probability vector with $pQ = 0$, and $D_{\text{KL}}(q(t) \| p) \rightarrow 0$ as $t \rightarrow \infty$; equivalently $q(t) \rightarrow p$ in any norm on \mathbb{R}^Ω . If Q is reducible, the total probability in each closed communicating class is conserved; the limit is obtained by rescaling p on each class to match those conserved masses, and $D_{\text{KL}}(q(t) \| p)$ relaxes to the corresponding constrained minimum rather than to zero.

Proof. Differentiating $D_{\text{KL}}(q(t) \| p) = \sum_i q_i(t) \log x_i(t)$ and using the master equation (23) gives $\frac{d}{dt} D_{\text{KL}} = \sum_i \dot{q}_i \log x_i + \sum_i \dot{q}_i = \sum_i \dot{q}_i \log x_i$ (since $\sum_i \dot{q}_i = 0$). Substituting $\dot{q}_j = \sum_i q_i Q_{ij}$ and reorganizing gives

$$\frac{d}{dt} D_{\text{KL}} = \sum_{i \neq j} q_i Q_{ij} \log x_j - \sum_{i \neq j} q_i Q_{ij} \log x_i,$$

where the second sum uses $\sum_j Q_{ij} = 0$. Detailed balance $p_i Q_{ij} = p_j Q_{ji}$ writes $q_i Q_{ij} = p_i Q_{ij} \cdot x_i = p_j Q_{ji} \cdot x_i$. Symmetrizing the indices in unordered pairs $\{i, j\}$ and writing $F_{ij} := p_i Q_{ij} = p_j Q_{ji}$ yields

$$\frac{d}{dt} D_{\text{KL}} = -\frac{1}{2} \sum_{i \neq j} F_{ij} (x_j - x_i) (\log x_j - \log x_i) \leq 0,$$

because $(x_j - x_i)(\log x_j - \log x_i) \geq 0$ for any pair of positive numbers and $F_{ij} \geq 0$. Convergence to the unique stationary law under irreducibility is the standard ergodic theorem for finite-state continuous-time Markov chains [31, Theorem 3.5.5]. \square

The contribution from each unordered pair $\{i, j\}$ vanishes iff $x_i(t) = x_j(t)$, so the equality conditions of (24) encode the fixed points of the dynamics. In the reducible case the equality is attained as soon as x_i is constant on each communicating class.

Combining this monotonicity with the algebraic identity of Proposition 6.1 gives the monotonicity of recognition free energy along the same dynamics, via a one-line differentiation.

Corollary 6.5 (Monotone recognition free energy under reversible Markov dynamics). *Fix $T_R > 0$ and let p be the Gibbs measure from Proposition 6.1. Under the hypotheses of Proposition 6.4, along any solution of (23) in the interior of the simplex,*

$$\frac{d}{dt} F_R(q(t)) \leq 0.$$

Proof. Differentiate the identity $F_R(q(t)) - F_R(p) = T_R D_{\text{KL}}(q(t)||p)$ from Proposition 6.1 in t , then apply Proposition 6.4 to bound the right-hand derivative. \square

Thus Proposition 6.1 and Corollary 6.2 are static identities on the simplex, while Proposition 6.4 supplies a standard reversible Markov dynamics for which $D_{\text{KL}}(\cdot||p)$ decreases. Corollary 6.5 then identifies the same decay as a monotone decrease of F_R . Together, these classical statements integrate the RCL-fixed cost vector $X_\omega = J(r_\omega)$ into the standard finite-state free-energy/relative-entropy chain [18, 19, 20, 11]. More general paths inherit the same conclusion only when $D_{\text{KL}}(q_t||p)$ is known to be nonincreasing; the physical content then comes from the chosen dynamics, not from the algebraic identity alone, and this kinematic statement should not be conflated with the Boltzmann H -theorem for an isolated Hamiltonian gas microstate.

7. Illustrative example: a three-state reversible system

We illustrate the results of Sections 2–6 on a three-state reversible Markov chain—equivalently, a chemical-reaction triangle $A \rightleftharpoons B \rightleftharpoons C \rightleftharpoons A$ at chemical equilibrium—building on the two-state setting recorded in Section 4.5. The example is constructed so that the input ratios r_ω have a transparent physical meaning: they are detailed-balance odds ratios of the reversible system, and the cost vector $X_\omega = J(r_\omega)$ is then determined by the closed form (5). The variational machinery of Section 5 produces the unique entropy-maximizing distribution at a chosen mean cost. We then check the free-energy–KL identity of Section 6 numerically, illustrate the squared-log bound of Proposition 3.1, and give a concrete reversible Markov trajectory illustrating the detailed-balance relaxation statement of Proposition 6.4.

7.1. Setup: input ratios from detailed balance

Consider a reversible three-state continuous-time Markov chain on $\Omega = \{1, 2, 3\}$ with strictly positive stationary law π and reference state 1 (the most populated species,

taken as $r_1 := 1$). Detailed balance $\pi_i W_{ij} = \pi_j W_{ji}$ identifies, for each non-reference state $\omega \in \{2, 3\}$, the dimensionless odds ratio

$$r_\omega := \frac{\pi_\omega}{\pi_1} = \frac{W_{1\omega}}{W_{\omega 1}},$$

exactly as in the two-state setting of Section 4.5. Once ratios are mapped to costs $X_\omega = J(r_\omega)$, Theorem 4.4 governs counting asymptotics for arbitrary real costs (hence generic irrational measured ratios); Corollary 4.5 records the classical exact-shell refinement in the rational regime. We take the specific values

$$r_2 := \frac{3-\sqrt{5}}{2} \approx 0.38197, \quad r_3 := 2 - \sqrt{3} \approx 0.26795,$$

which are the unique solutions of $J(r_2) = \frac{1}{2}$ and $J(r_3) = 1$ in $(0, 1)$ under (5), chosen so that the cost vector takes the simple form $X = (0, \frac{1}{2}, 1)$; in closed form, $r_2 = \varphi^{-2}$ with $\varphi = (1 + \sqrt{5})/2$ the golden ratio, and $r_3 = \tan(\pi/12) = 2 - \sqrt{3}$. These choices make the costs rational, so Corollary 4.5 (exact integer shell along denominator-exhaustive sequences) applies directly in this subsection. A worked illustration with $r_2 = e^{-1}$, $r_3 = e^{-3/2}$, for which the exact shell is empty at every N but Theorem 4.4 still applies, is given in Section 8.8. The chain's own stationary law follows from $\pi_\omega \propto r_\omega$ and normalization:

$$\pi = \left(\frac{1}{1 + r_2 + r_3}, \frac{r_2}{1 + r_2 + r_3}, \frac{r_3}{1 + r_2 + r_3} \right) \approx (0.60609, 0.23151, 0.16240).$$

The ratios r_ω thus encode physically meaningful equilibrium odds: state 1 (the reference) is the most populated, with $r_2 \approx 0.38$ giving roughly 38% as much population in state 2 relative to 1, and $r_3 \approx 0.27$ giving about 27% as much in state 3. In a reaction-network reading, r_2 and r_3 are equilibrium constants for the reactions $1 \rightarrow 2$ and $1 \rightarrow 3$, both unfavorable (equilibrium constants below unity); the corresponding affinities $\Delta\mu_\omega/k_B T = -\log r_\omega$, evaluated at the convention $k_B T = 1$ (so that $\Delta\mu_\omega$ is reported in dimensionless units of $k_B T$), are $\Delta\mu_2 \approx 0.962$ and $\Delta\mu_3 \approx 1.317$.

7.2. Gibbs distribution at fixed mean cost

With the cost vector $X = (0, \frac{1}{2}, 1)$ fixed by the input ratios, fix a constraint value $E_{\text{mc}} = 0.45$. This value lies strictly between $\min X = 0$ and the uniform mean $\langle X \rangle_{p^{(0)}} = \frac{1}{3}(0 + \frac{1}{2} + 1) = \frac{1}{2}$, so Proposition 5.2 guarantees a unique solution $\beta > 0$ to the moment-matching equation

$$\frac{0 \cdot e^0 + (1/2)e^{-\beta/2} + 1 \cdot e^{-\beta}}{1 + e^{-\beta/2} + e^{-\beta}} = E_{\text{mc}}, \quad (25)$$

which is $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$ with $X = (0, \frac{1}{2}, 1)$ written out. Numerical solution gives

$$\beta \approx 0.30113, \quad T_{\text{R}} = \beta^{-1} \approx 3.3208, \quad Z \approx 2.6002,$$

and Gibbs weights

$$p \approx (0.38459, 0.33083, 0.28459),$$

monotonically decreasing in the cost: the reference state has the largest weight, and the highest-cost state the smallest. For comparison, the uniform reference distribution $q = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ corresponds to $\beta = 0$ and the higher mean cost $\langle X \rangle_q = \frac{1}{2}$.

The Gibbs distribution p is the maximum-entropy distribution among all probability laws on Ω with the prescribed mean cost $E_{\text{mc}} = 0.45$. It need not coincide with the chain's own stationary distribution $\pi \approx (0.6061, 0.2315, 0.1624)$ from the previous subsection: π has its own realized mean cost $\langle X \rangle_\pi = (r_2/2 + r_3)/(1 + r_2 + r_3) \approx 0.2782$, distinct from E_{mc} . The two distributions coincide only when the chain's stationary distribution is itself a Gibbs measure for the cost vector X at the matching inverse temperature: $\pi = p^{(\beta)}$ for some $\beta > 0$ if and only if the chosen constraint value E_{mc} equals $\langle X \rangle_\pi$, and the matching β is the unique root from Proposition 5.2. The Gibbs distribution of Section 5 is a thermodynamic object selected by entropy maximization at fixed mean cost, while the chain's stationary law is a kinematic object determined by its rate matrix; the two come from compatible but distinct constructions.

7.3. Numerical values of the divergences

With p and q specified, the functionals of Sections 3–6 take the closed-form numerical values shown in Table 2. The free-energy–KL identity of Proposition 6.1 is an algebraic equality and the agreement of the rows $F_{\text{R}}(q) - F_{\text{R}}(p)$ and $T_{\text{R}} D_{\text{KL}}(q||p)$ in Table 2 only documents that the numerical implementation is consistent with the closed-form identity. The interesting comparison is between D_J and the squared-log surrogate of Lemma 2.1: the inequality $D_J(q||p) \geq \frac{1}{2} \sum_{\omega} p_{\omega} (\log(q_{\omega}/p_{\omega}))^2$ from Proposition 3.1 (iii) is satisfied with the small positive margin $0.00751 - 0.00750 \approx 1 \times 10^{-5}$ visible in the last two rows, as expected near coincidence. The next subsection turns to a genuinely dynamical illustration of the detailed-balance Lyapunov statement.

7.4. Detailed-balance relaxation trajectory

To put Proposition 6.4 on this example, we equip the three states with an irreducible continuous-time Markov chain that has p as its unique stationary law and is reversible with respect to it. The Metropolis-on-a-graph rates [23],

$$Q_{ij} = \min\left(1, \frac{p_j}{p_i}\right) \quad (i \neq j), \quad Q_{ii} = -\sum_{j \neq i} Q_{ij},$$

satisfy detailed balance $p_i Q_{ij} = p_j Q_{ji}$ by construction, with row sums zero. With the Gibbs reference $p \approx (0.385, 0.331, 0.285)$ from Section 7.3, the off-diagonal block evaluates to

$$Q_{12} = 0.860, \quad Q_{13} = 0.740, \quad Q_{21} = 1.000, \quad Q_{23} = 0.860, \quad Q_{31} = Q_{32} = 1.000.$$

Starting from the uniform initial law $q(0) = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ and integrating $\dot{q} = qQ$ via $q(t) = q(0) \exp(tQ)$ gives the trajectory in Table 3.

Quantity	Value (rounded)
$H(p)$	1.0911
$H(q)$	$\log 3 \approx 1.0986$
$D_{\text{KL}}(q p)$	0.00754
$F_{\text{R}}(p)$	-3.1733
$F_{\text{R}}(q)$	-3.1483
$F_{\text{R}}(q) - F_{\text{R}}(p) = T_{\text{R}} D_{\text{KL}}(q p)$	0.02505
$D_J(q p)$	0.00751
$\frac{1}{2} \sum_{\omega} p_{\omega} (\log(q_{\omega}/p_{\omega}))^2$	0.00750

p denotes the Gibbs measure at mean cost $E_{\text{mc}} = 0.45$;
 $q = (\frac{1}{3}, \frac{1}{3}, \frac{1}{3})$ is the uniform reference distribution.

Table 2. Three-state example with $E_{\text{mc}} = 0.45$. The row $F_{\text{R}}(q) - F_{\text{R}}(p) = T_{\text{R}} D_{\text{KL}}(q||p)$ records the algebraic identity of Proposition 6.1; the last two rows illustrate the squared-log lower bound of Proposition 3.1 (iii).

t	$q_1(t)$	$q_2(t)$	$q_3(t)$	$D_{\text{KL}}(q(t) p)$	$F_{\text{R}}(q(t))$
0.0	0.3333	0.3333	0.3333	7.54×10^{-3}	-3.14830
0.5	0.3706	0.3323	0.2970	5.29×10^{-4}	-3.17159
1.0	0.3808	0.3314	0.2878	3.74×10^{-5}	-3.17322
2.0	0.3843	0.3309	0.2848	1.94×10^{-7}	-3.17335
5.0	0.3846	0.3308	0.2846	$< 10^{-12}$	-3.17335

Table 3. Relaxation under the Metropolis chain reversible with respect to p , started from the uniform law. Both $D_{\text{KL}}(q(t)||p)$ and $F_{\text{R}}(q(t))$ are monotonically decreasing in t , as predicted by Proposition 6.4 and Corollary 6.5; $q(t) \rightarrow p$ exponentially, and $F_{\text{R}}(q(t))$ approaches the equilibrium value $F_{\text{R}}(p) \approx -3.1733$.

The trajectory is a direct numerical realization of Proposition 6.4: D_{KL} falls from 7.54×10^{-3} to below 10^{-12} over five (arbitrary) Markov-time units, and $F_{\text{R}}(q(t)) - F_{\text{R}}(p) = T_{\text{R}} D_{\text{KL}}(q(t)||p)$ tracks the same exponential decay rescaled by $T_{\text{R}} = 1/\beta \approx 3.32$. The leading relaxation rate is set by the spectral gap of Q , and it is the standard exponential approach to canonical equilibrium for the chosen rates; what changes from one model to another is only the cost vector X that enters p via the variational principle.

7.5. Comparison with the Tsallis q -exponential alternative

The numerical agreement above checks an algebraic identity. To turn the example into a non-trivial comparison, we juxtapose the RCL Gibbs distribution with the Tsallis q -exponential obtained under an alternative entropy functional on the same cost data. To avoid clashing with the symbol q already used for probability distributions, we write q_T for the Tsallis deformation index throughout this subsection.

The Tsallis program [9, 10] replaces Shannon entropy by the q_T -deformed entropy

$$S_{q_T}(p) := \frac{1}{q_T - 1} \left(1 - \sum_{\omega \in \Omega} p_{\omega}^{q_T} \right), \quad q_T > 0,$$

which reduces to H as $q_T \rightarrow 1$. Maximizing S_{q_T} at fixed $\langle X \rangle_p = E_{\text{mc}}$ on the simplex (constraint-1 formulation; see [9, 10]) yields

$$p_{\omega}^{(\beta_{q_T})} \propto [1 - (1 - q_T)\beta_{q_T}X_{\omega}]_+^{1/(1-q_T)},$$

with β_{q_T} adjusted so that $\sum_{\omega} p_{\omega}^{(\beta_{q_T})} X_{\omega} = E_{\text{mc}}$. For $q_T = 1$ this reduces to the Gibbs form; for $q_T \neq 1$ the tails are power-law rather than exponential. Table 4 fits β_{q_T} to the same constraint $E_{\text{mc}} = 0.45$ as Table 2 and reports the ℓ^1 distance to the RCL Gibbs measure $p^{\text{RCL}} \approx (0.385, 0.331, 0.285)$.

Scope of the comparison. This juxtaposition is a controlled deformation of the Gibbs branch within the fixed RCL cost vector $X = (0, \frac{1}{2}, 1)$ and the simple-mean (constraint-1) formulation used throughout this paper. It is not a comparison against the full Tsallis statistical-mechanical framework, where the natural cost choice and the escort-mean (constraint-3) formulation may differ from the present setup; conclusions drawn from Table 4 should not be read as general statements about Tsallis non-extensive statistical mechanics. Its purpose is only to illustrate that, on the same finite cost data and mean-cost constraint, deforming the entropy functional within the q_T -family yields distributions that differ from the Gibbs branch in a controlled way and become indistinguishable as $q_T \rightarrow 1$.

Distribution	Tsallis index q_T	β_{q_T}	p_1	p_2	p_3	$\ \cdot - p^{\text{RCL}}\ _1$
RCL Gibbs	1.00	0.3011	0.3846	0.3308	0.2846	0
Tsallis q_T -exp	0.10	0.2646	0.3835	0.3331	0.2835	0.0045
Tsallis q_T -exp	0.50	0.2801	0.3840	0.3321	0.2840	0.0025
Tsallis q_T -exp	1.50	0.3243	0.3852	0.3296	0.2852	0.0025
Tsallis q_T -exp	2.50	0.3780	0.3864	0.3272	0.2864	0.0073

Table 4. Three-state example with $X = (0, \frac{1}{2}, 1)$ and $E_{\text{mc}} = 0.45$. Tsallis q_T -exponential distributions match the mean-cost constraint at different shape parameters β_{q_T} ; the resulting weights differ from the RCL Gibbs measure by an ℓ^1 amount that increases as the deformation moves away from the Gibbs case $q_T = 1$. The gap is small at $E_{\text{mc}} = 0.45$ because the constraint sits close to the uniform-cost value 0.5, where the high-temperature limit flattens all q_T -exponential families to the uniform measure.

At the matched constraint $E_{\text{mc}} = 0.45$, the Tsallis entropy $S_{q_T}(p^{(\beta_{q_T})})$ and the Helmholtz analogue $F^{(q_T)} = \langle X \rangle_p - T_{q_T} S_{q_T}(p)$ with $T_{q_T} := 1/\beta_{q_T}$ evaluate to $S_{q_T=0.5} = 1.4578$ and $F^{(0.5)} \approx -4.7543$; $S_{q_T=1} = H(p^{\text{RCL}}) = 1.0911$ and $F_{\text{R}}(p^{\text{RCL}}) \approx -3.1737$; $S_{q_T=1.5} = 0.8388$ and $F^{(1.5)} \approx -2.1365$; and $S_{q_T=2.5} = 0.5347$ and $F^{(2.5)} \approx -0.9646$.

These values are computed analytically from the Table 4 weights and the fitted β_{q_T} entries.

For $q_T < 1$ the population shifts slightly toward the median-cost state and away from the extremes within this fixed-cost-and-constraint setup; for $q_T > 1$ the opposite occurs. The separation grows as E_{mc} moves away from the uniform value: at $E_{\text{mc}} = 0.20$, $\|p_{q_T=0.5}^{\text{Tsallis}} - p^{\text{RCL}}\|_1 \approx 0.087$ and $\|p_{q_T=2.5}^{\text{Tsallis}} - p^{\text{RCL}}\|_1 \approx 0.127$, an order of magnitude larger than at $E_{\text{mc}} = 0.45$. Within this deformation, the same mean cost is therefore compatible with a one-parameter family of distributions; concretely, the observables that resolve the families are those sensitive to more than the mean of X , including the variance $\text{Var}_p(X)$, the high-cost tail probability p_3 , and any moment $\langle X^k \rangle_p$ for $k \geq 2$. The comparison itself is not an empirical validation of the RCL, and it is not a comparison against the full Tsallis non-extensive statistical-mechanical framework: it only demonstrates that, within a fixed cost vector and simple-mean constraint, deforming the entropy functional within the q_T -family yields finite-state distributions distinguishable from the Gibbs branch at the fluctuation level rather than at the mean.

7.6. Sample-size requirements at fixed RCL ground truth (power calculation)

This subsection is a power calculation, not a falsifiability test of the RCL: the RCL Gibbs law p^{RCL} is taken as the generating distribution, and each alternative framework is asked how many effectively independent observations would be needed to resolve its predicted tail weight from that of p^{RCL} at nominal 2σ on a single coordinate. An actual falsifiability test would require an experimental three-state system with independently measured rate ratios and stationary populations.

Throughout this subsection and in Table 5, the quantity $N_{2\sigma}$ denotes the *effective* number of independent Bernoulli trials equivalent, for inferring the frequency \hat{p}_3 , to the Wald two-sided $z = 2$ sample size after correcting for temporal autocorrelation in Markovian trajectories (thinning, block averaging, integrated autocorrelation time, or related standard procedures).

An ℓ^1 distance between two candidate distributions is a useful summary statistic, but it does not directly tell an experimenter whether the two distributions can be distinguished at a specified confidence level. We therefore convert the model gap into a sample-size requirement on the most informative single coordinate.

The Tsallis comparison of Section 7.5 reports an ℓ^1 distance, but does not specify which physical observable would discriminate between the alternatives or how many samples would be needed to do so. The natural finite-state observable is the population frequency \hat{p}_ω on each state, accumulated from N independent observations. Among the three frequencies, the highest-cost weight p_3 is the most directly informative, since it is the most sensitive to the cost choice in the Gibbs construction.

To put cost frameworks on equal footing, we adopt the following protocol. Take the RCL Gibbs distribution p^{RCL} at $E_{\text{mc}} = 0.45$ as the generating alternative in a model-discrimination power calculation (Section 7.3). Each candidate cost framework is

asked to explain that physical state by fitting its inverse temperature to the empirical mean of its own cost computed at p^{RCL} . The four candidates on the same ratio data $(r_1, r_2, r_3) = (1, \varphi^{-2}, \tan(\pi/12))$ are: the RCL cost $X_\omega^{\text{RCL}} = J(r_\omega) = (0, \frac{1}{2}, 1)$ (Section 7); the squared-log surrogate $X_\omega^{\text{SQ}} = \frac{1}{2}(\log r_\omega)^2 \approx (0, 0.463, 0.867)$ from Lemma 2.1; the affinity-as-energy assignment $X_\omega^{\text{A}} = -\log r_\omega \approx (0, 0.962, 1.317)$, which treats the reference-oriented log-ratio directly as the energy-like coordinate (the canonical-ensemble version of the stochastic-thermodynamic comparison in Section 8.3); and the Tsallis q_T -exponential at the same RCL cost vector and constraint value (Table 4).

Table 5 reports the predicted p_3 for each framework together with $N_{2\sigma}$ as defined in the table caption, required to distinguish that prediction from $p_3^{\text{RCL}} = 0.2846$ at the 2σ level on a single coordinate of the empirical type. This estimate is a deliberately simple power calculation: it uses a prespecified single-coordinate statistic and the binomial standard error $\sigma_{\hat{p}_3}^2 = p_3(1 - p_3)/N$ only after N has been reduced to an effective independent-sample count as in the opening paragraph. A multi-coordinate goodness-of-fit test would require the corresponding likelihood-ratio or multiple-comparison calibration. The affinity-as-energy alternative is the most readily distinguishable: at $E_{\text{mc}} = 0.45$ it predicts $p_3 \approx 0.295$ versus the RCL value 0.285, a binomial-distinguishable difference at $N \approx 8 \times 10^3$ —in principle accessible in systems that provide on the order of 10^4 effectively independent state observations after autocorrelation correction. The squared-log surrogate and the Tsallis $q_T = 2.5$ case require $N \sim 10^5$ – 10^6 , which is more demanding and would be realistic only in high-throughput or highly repeatable assays. The Tsallis case at $q_T = 0.5$, very close to the Gibbs branch $q_T = 1$, is essentially indistinguishable.

Framework	Cost vector X_ω	Fitted β	p_3	Δp_3	$N_{2\sigma}$
RCL Gibbs	(0, 0.500, 1.000)	0.301	0.2846	0	—
Squared-log surrogate	(0, 0.463, 0.867)	0.346	0.2858	+0.0012	5.7×10^5
Affinity-as-energy	(0, 0.962, 1.317)	0.210	0.2945	+0.0099	8.3×10^3
Tsallis $q_T = 0.5$	(0, 0.500, 1.000)	0.280	0.2840	-0.0006	2.1×10^6
Tsallis $q_T = 2.5$	(0, 0.500, 1.000)	0.378	0.2864	+0.0018	2.4×10^5

Table 5. Three-state example at $E_{\text{mc}} = 0.45$. Each candidate framework fits its inverse temperature β to match the mean of its own cost computed at p^{RCL} , and predicts the heavy-state weight p_3 . The column $N_{2\sigma}$ is the Wald two-sided $z = 2$ binomial sample-size formula $N_{2\sigma} \approx 4p_3(1 - p_3)/(\Delta p_3)^2$, i.e. the number of *effectively independent* observations of state 3 needed to resolve the shift Δp_3 from $p_3^{\text{RCL}} = 0.2846$ at nominal 2σ on \hat{p}_3 alone, after reducing raw trajectory length to an equivalent independent-sample count when data are Markov-correlated (cf. the opening of Section 7.6).

The discriminating power increases sharply as E_{mc} moves away from the high-temperature limit $E_{\text{mc}} \rightarrow \frac{1}{2}$. At $E_{\text{mc}} = 0.20$ (recognition temperature $T_{\text{R}} \approx 0.47$, deeper into the cost-penalized regime), the same protocol gives $p_3^{\text{RCL}} \approx 0.082$ and predicted differences $\Delta p_3 \approx 0.035$ for the affinity-as-energy alternative, 0.032 for Tsallis $q_T = 2.5$, and 0.022 for Tsallis $q_T = 0.5$, requiring binomial samples of only $N_{2\sigma} \approx 250, 300$, and 640 respectively—within reach of moderate-throughput dwell-time histograms or

equilibrium-population assays in favorable systems [32, 33]. The squared-log surrogate, which agrees with the RCL cost only to second order at $r = 1$, becomes distinguishable at $N_{2\sigma} \approx 1.3 \times 10^4$ in this regime.

The qualitative content is therefore that the present finite-state framework makes quantitatively distinct predictions from each of the natural alternatives—squared-log, affinity-as-energy, and Tsallis—on the same input ratios, and that the most directly competing framework (affinity-as-energy, which one would naturally adopt by treating the dimensionless edge affinity $\log r_\omega$ as an energy) is distinguishable at sample sizes that are, in principle, accessible after autocorrelation correction over the illustrated parameter range, becoming sharply distinguishable away from the uniform-cost limit. The discriminating observable need not be sophisticated: a single coordinate of the empirical type vector suffices. These estimates measure how distinguishable the RCL prediction is from each alternative at fixed mean cost, conditional on the RCL Gibbs law being the generating distribution. They are not a falsifiability test of the RCL itself; an actual test would require an experimental three-state system with independently measured rate ratios and stationary populations, so that the RCL, affinity-as-energy, and Tsallis predictions could each be compared against measured data.

7.7. Dependence on the constraint

Varying E_{mc} within its admissible range $(0, \frac{1}{2})$ while keeping the cost vector $X = (0, \frac{1}{2}, 1)$ fixed gives a more global view. As E_{mc} increases from 0 toward the uniform mean $\frac{1}{2}$, the canonical parameter β from Proposition 5.2 decreases from $+\infty$ to 0^+ , the recognition temperature $T_{\text{R}} = 1/\beta$ rises from 0 to $+\infty$, and the Gibbs measure deforms continuously from the point mass on the zero-cost outcome to the uniform measure. Table 6 lists the canonical parameters at three sample points approaching the midpoint $E_{\text{mc}} = \frac{1}{2}$ from below; the rapid divergence of T_{R} and the flattening of the Gibbs weights are visible from the numerical entries.

Figure 2 displays this dependence graphically. The weight p_1 on the zero-cost outcome decays monotonically as E_{mc} increases, while p_3 on the highest-cost outcome rises monotonically; p_2 traces a milder curve between the two. All three curves converge to $\frac{1}{3}$ as $E_{\text{mc}} \rightarrow \frac{1}{2}$, the high-temperature limit in which the cost contribution to the Gibbs weights becomes negligible. This deformation is the finite-state shadow of the standard high-temperature limit of the canonical ensemble.

$\langle X \rangle$ (target)	β	$T_{\text{R}} = 1/\beta$	Z	p_1 (rounded)
0.35	0.932	1.073	2.021	0.495
0.45	0.301	3.321	2.600	0.385
0.48	0.120	8.328	2.829	0.354

Table 6. Same three-state costs $X = (0, \frac{1}{2}, 1)$ as Table 2; canonical parameters as the target mean cost $\langle X \rangle$ (the constraint value denoted E_{mc} elsewhere) varies (numerical root of $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$).




Figure 2. Gibbs weights p_k as a function of the target mean cost $\langle X \rangle$ for the three-state example with costs $X = (0, 0.5, 1)$, plotted over the positive-temperature range $E_{\text{mc}} \in (0, \frac{1}{2})$ in which $\beta > 0$ (Proposition 5.2); E_{mc} denotes the same constraint as the column $\langle X \rangle$ (target) in Table 6. The dotted vertical line marks the high-temperature limit $E_{\text{mc}} \rightarrow \frac{1}{2}$ at which $\beta \rightarrow 0^+$ and the distribution approaches the uniform measure.

8. Discussion

The derivation keeps three steps explicit and separate: fix the scalar cost from the RCL, recover Shannon entropy from multinomial counting, and apply convex optimization to obtain Gibbs weights. Once X_ω is fixed, Section 5 reproduces the maximum-entropy characterization of the Gibbs measure and Section 6 the free-energy–Kullback–Leibler identity for the canonical exponential family on a finite space.

8.1. Logical summary of the derivation

Table 7 lists the nine steps of the derivation in order; each row is a consequence of those above together with a single new ingredient (calibration, counting, or convex analysis).

8.2. Relation to Jaynes and standard statistical mechanics

In Jaynes’s maximum-entropy program [4, 5, 6], a scalar cost—typically a Hamiltonian or phenomenological energy—is chosen, its mean is fixed on the simplex, and Shannon entropy is maximized, yielding the Gibbs weights and an inverse temperature as Lagrange multiplier. The remaining modeling freedom lies in the choice of cost functional and constraint value; multinomial counting, Stirling entropy, and the convex duality identities are shared with the present work.

The present work differs in one place: the cost is selected upstream by the RCL together with (A2), rather than chosen as an external Hamiltonian. Once dimensionless input ratios $\{r_\omega\}$ are specified, $X_\omega = J(r_\omega)$ has the closed form (5). Modeling freedom

Step	Result	Reference
1	(A1)–(A3) $\Rightarrow J(x) = \frac{1}{2}(x + x^{-1}) - 1$	Theorem 1.1
2	$J(x) \geq (\log x)^2/2$	Lemma 2.1
3	D_J properties; half-Neyman χ^2 form	Proposition 3.1
4	Shannon entropy from constrained multinomials	Theorem 4.4; Corollary 4.5; non-asymptotic bound, Proposition 4.1; multi-constraint soft shell, Proposition 8.1
5	Existence and uniqueness of β ($T_R = 1/\beta$ for $\beta > 0$)	Lemma 5.1, Proposition 5.2
6	Finite-state Gibbs variational principle	Proposition 5.3
7	$F_R(q) - F_R(p) = T_R D_{\text{KL}}(q p)$	Proposition 6.1
8	Gibbs minimizes F_R on the simplex	Corollary 6.2
9	D_{KL} nonincreasing under detailed balance	Proposition 6.4, Corollary 6.5

Table 7. Step-by-step derivation chain. Origins of each result and the contribution categories (N1)–(N3) are stated in Section 1.3.

is relocated from the choice of cost functional to the choice of ratio assignment, not eliminated: specifying $\{r_\omega\}$ is itself a modeling commitment of comparable difficulty to specifying a Hamiltonian, and the Tsallis comparison of Section 7.5 shows that even with $\{r_\omega\}$ fixed there is residual freedom in choosing the composition law.

The Shore–Johnson axiomatization [7, 8] is complementary. Shore and Johnson fix Shannon entropy as the unique self-consistent inference rule under linear constraints; here Shannon entropy is taken as given, emerging from multinomial counting in Section 4, and the focus is on which cost the entropy is maximized against. In one slogan: Shore–Johnson constrains the entropy; the RCL constrains the cost. Caticha’s entropic inference [34, 8] extends the Shore–Johnson program into a broader epistemic setting on which the RCL is silent. A natural continuation, which we leave open, is to ask whether the Shore–Johnson axioms restated in ratio coordinates would force, or merely permit, the d’Alembert composition law (2); we return to this in Section 9.

From the large-deviation perspective, Sanov’s theorem and its conditional formulations describe how empirical measures concentrate under linear constraints in the large- N limit [29, 22, 21], with the Gibbs measure as the unique minimizer of the rate function. Theorem 4.4 is the finite- K method-of-types instance of this picture [20], with rate of convergence controlled by Proposition 4.1; see [11, 18] for the broader framework. The information-geometric viewpoint of Amari [35, 36] describes $\{p^{(\beta)}\}_{\beta \in \mathbb{R}}$ as a one-dimensional exponential family with Fisher metric induced by $\log Z$ and Proposition 6.1 as the Legendre–Fenchel duality between natural and expectation parameters; this geometry is the natural language for extensions to higher-dimensional sufficient statistics, but is not pursued here.

Neighboring literatures on the cosh-kernel and on multi-variable functional equations. The closed cost $J(e^t) = \cosh t - 1$ is a hyperbolic-cosine kernel in the log-ratio coordinate $t = \log r$, a form that appears in several neighboring information-theoretic and information-geometric literatures. The Itakura–Saito divergence [37] and its symmetrized

variants on the positive cone $\mathbb{R}_{>0}$ are ratio-based contrasts—built directly from $r = q/p$ —in which the symmetric (even) component is of the same hyperbolic-cosine type as J , while the antisymmetric (odd) component is the log-ratio that also appears in the affinity $\log r$ of Section 8.3; in the symmetric case both Itakura–Saito and the present J vanish quadratically at $r = 1$ and grow superquadratically in $|\log r|$. On positive-definite matrices, the symmetrized Bregman or log-determinant Stein divergences [38, 39] exhibit a related cosh-of-log-eigenvalue structure as the natural even part of the asymmetric log-determinant divergence. The Gibbs construction here recovers the one-dimensional exponential family with cumulant generator $\log Z(\beta)$; whether the RCL Gibbs measure carries an additional information-geometric interpretation beyond this family—for example a non-trivial dually-flat structure on the positive cone induced by D_J —is not addressed in this paper. For the multi-variable functional-equation theory used in Theorem 1.1 (the d’Alembert equation on \mathbb{R}^2 in log-coordinates) the standard reference beyond Aczél [14] and Kuczma–Choczewski–Ger [15] is Aczél and Dhombres [40].

8.3. Connection to stochastic thermodynamics: $J(r)$ as one symmetric summary of the edge data

The choice to enter the ratios $\{r_\omega\}$ through the cost $X_\omega = J(r_\omega)$, rather than through the affinity $\log r_\omega$, has a direct interpretation in the language of stochastic thermodynamics [32, 41]. Consider a reversible continuous-time Markov chain on Ω with positive transition rates W_{ij} ($i \neq j$) satisfying detailed balance $\pi_i W_{ij} = \pi_j W_{ji}$ with respect to a strictly positive stationary law π [23]. Each edge $i \leftrightarrow j$ then carries a single dimensionless ratio

$$r_{ij} := W_{ij}/W_{ji} = \pi_j/\pi_i, \quad r_{ji} = r_{ij}^{-1}.$$

This scalar admits two natural summaries on the same edge data. The rate-ratio contribution to the stochastic-thermodynamic force is

$$A_{ij}^{\text{rate}} := \log r_{ij} = -A_{ji}^{\text{rate}}, \quad (26)$$

which is antisymmetric under edge reversal. At a nonequilibrium distribution q , the corresponding Schnakenberg force is

$$\mathcal{A}_{ij}(q) := \log \frac{q_i W_{ij}}{q_j W_{ji}} = \log \frac{q_i}{q_j} + A_{ij}^{\text{rate}}.$$

The edge entropy-production contribution is

$$(q_i W_{ij} - q_j W_{ji}) \mathcal{A}_{ij}(q) \geq 0,$$

and the full force $\mathcal{A}_{ij}(q)$, not the bare rate-ratio contribution A_{ij}^{rate} , vanishes at detailed-balance equilibrium $q = \pi$ [19, 32, 41]. One symmetric summary of the same edge-ratio data is the RCL cost

$$J(r_{ij}) = \cosh A_{ij}^{\text{rate}} - 1 = J(r_{ji}) \quad (27)$$

(Any even function f with $f(0) = 0$ furnishes an alternative symmetric summary of the same edge ratio, for instance via $f(\log r_{ij})$ or $f(|A_{ij}^{\text{rate}}|)$.) It is one natural choice among them, selected by the RCL axioms (S1)–(S3) and the degree-two closure of Section 1.2: it is even under edge reversal, depends only on $|A_{ij}^{\text{rate}}|$, and vanishes precisely when the bare edge ratio is unity, $r_{ij} = 1$, i.e. when the two oriented rates are equal before weighting by the stationary distribution. This is distinct from detailed-balance equilibrium, for which the full force $\mathcal{A}_{ij}(q)$ vanishes at $q = \pi$ even when $r_{ij} \neq 1$. Canonicity is not claimed, and proving that J is the unique such functional compatible with detailed-balance entropy-production inequalities or current-fluctuation symmetries would require additional axioms not stated here.

This parity decomposition gives the cost $X_\omega = J(r_\omega)$ a recognizable physical role: it is a symmetric, equilibrium-oriented summary of the same rate-ratio data whose antisymmetric component enters stochastic thermodynamics through forces and currents. The Gibbs construction of Sections 5–6 uses only this symmetric content J , consistent with its restriction to finite-state equilibrium canonical-ensemble objects; the decomposition itself, however, is not unique (any even function of $\log r$ with $f(0) = 0$ furnishes a symmetric summary), and the same latitude applies when reading J against the stochastic-thermodynamic literature. Cycle affinities, broken-detailed-balance currents, and trajectory-level fluctuation relations [42, 33, 41] would require retaining the antisymmetric/current component as additional structure, in line with the information-geometric thermodynamic interpretation of Ito [43] and the thermodynamic uncertainty relations bounding precision against dissipation [44]. We do not pursue the corresponding non-equilibrium extension here; the present finite-state framework keeps only the symmetric content of each edge and is therefore restricted to detailed-balance equilibrium statistics.

8.4. Where the input ratios come from in practice

The framework treats $\{r_\omega\}$ as primary input. Three classes of finite-state systems supply such ratios directly, without an intervening Hamiltonian:

- (i) Reversible chemical equilibria. For a closed reaction network at thermodynamic equilibrium, equilibrium constants $K_{\text{eq}} = [\text{B}]/[\text{A}]$ are dimensionless ratios accessible from concentration measurements (spectroscopic absorbance ratios, mass spectrometry abundance ratios, NMR integration). These ratios are the natural inputs to a finite-state ratio-cost model whose states are reactant/product species and whose Gibbs distribution gives a canonical finite-state population model at the prescribed mean cost.
- (ii) Reversible single-molecule kinetics. Detailed-balance ratios W_{ab}/W_{ba} between transition rates of a reversible discrete-state Markov chain on conformations or binding states can be estimated directly from dwell-time histograms in single-molecule fluorescence or atomic-force experiments, even when the underlying free-energy profile along the chosen reaction coordinate is not directly accessible.

- (iii) Discrete-state inference from log-odds. In binary or multi-class classification settings where one observes posterior odds ratios across categories without an explicit log-likelihood, the ratios r_ω are the experimentally available primitives, and (20) produces a finite-state distribution whose mean cost matches an empirical constraint.

In each case the operative quantity is a ratio, not a Hamiltonian, and the relocation of modeling freedom from a cost functional to $\{r_\omega\}$ together with the d'Alembert composition law is a step toward more transparent inference rather than a thermodynamic re-derivation. The framework does not, on its own, predict the ratios; it converts them, once specified, into a canonical exponential family.

8.5. Other composition laws and other axiomatic frameworks

The motivation of (2) from (S1)–(S3) in Section 1.2 is not a uniqueness theorem. Several alternative composition laws on the ratios r are also consistent with the multiplicative algebra of detailed balance and equilibrium constants:

- (C1) Multiplicative Cauchy: $J(xy) = J(x)J(y)$ on $\mathbb{R}_{>0}$. With $J(1) = 0$ this forces $J \equiv 0$; the $J(1) \neq 0$ branches yield $J(x) = x^\lambda$ but lose (S2). The corresponding maximum-entropy distribution is a power-law family.
- (C2) Additive Cauchy in log-coordinates: $G(u+v) = G(u) + G(v)$ for $G(u) = J(e^u)$, with continuous solutions $G(u) = cu$, hence $J(x) = c \log x$. Violates the reciprocity condition (S1) for all $c \neq 0$, since then $J(x^{-1}) = -J(x)$, not $J(x)$.
- (C3) Squared-log surrogate (additive d'Alembert): $G(u+v) + G(u-v) = 2G(u) + 2G(v)$, with continuous solutions $J(x) = c(\log x)^2$; the choice $c = \frac{1}{2}$ is the surrogate of Lemma 2.1. Compatible with (S1)–(S2) but strictly weaker far from equilibrium than (5).
- (C4) Tsallis q -deformation: replaces J by $J_q(x) = (x^{1-q} - 1)/(1 - q)$ and H by the q -entropy S_q , yielding q -exponential distributions [9, 10]; numerical comparison in Section 7.5.
- (C5) Naudts ϕ -exponentials: replaces \exp by a deformed exponential ϕ^{-1} , with Tsallis as the special case $\phi(x) = x^q$ [10, 45].

Table 8 collects these alternatives, together with the one-parameter bilinear degree-two template in its second row: normalization $\alpha = 1$ recovers the RCL (2), while $\alpha = 0$ collapses to the additive (C3) branch and is excluded by the nonzero bilinear coupling adopted in Section 1.2. The RCL is one option among them, and its specific status is the following: under (S1)–(S3) and the additional normalized degree-two d'Alembert closure of Section 1.2, $J(x) = \frac{1}{2}(x + x^{-1}) - 1$ is the unique non-trivial continuous solution. Drop that closure and the alternatives (C1)–(C5) become admissible. Drop (S1) and the Tsallis and additive-Cauchy options become admissible. The RCL is therefore privileged only relative to the joint adoption of (S1)–(S3) plus the normalized degree-two d'Alembert closure; whether that joint adoption is the right modeling commitment is the substantive

question on which the entire framework rests, and the present paper does not attempt to settle it.

Composition law	Cost J on $\mathbb{R}_{>0}$	Symmetric in $x \leftrightarrow x^{-1}$?	MaxEnt distribution
RCL (2)	$\frac{1}{2}(x+x^{-1})-1$ (cosh $\log x - 1$)	yes	Gibbs $\propto e^{-\beta J}$
Bilinear degree-two (unnormalized α)	$J(xy) + J(x/y) = 2\alpha J(x)J(y) + 2J(x) + 2J(y)$; $\alpha = 1$ is (2); $\alpha = 0$ is the Cauchy/(C3) degeneracy excluded by nonzero bilinear coupling	yes (nontrivial branch)	Gibbs $\propto e^{-\beta J}$ (same class as RCL under $J \mapsto \alpha J, \beta \mapsto \beta/\alpha$)
Squared-log (C3)	$\frac{1}{2}(\log x)^2$	yes	log-Gaussian $\propto e^{-\beta(\log x)^2/2}$
Multiplicative Cauchy (C1)	x^λ (loses (S2))	no	power-law family
Additive Cauchy (C2)	$c \log x$ (loses (S1))	no	geometric $\propto x^{-\beta c}$
Tsallis deformation [9]	q - $J_q(x) = (x^{1-q} - 1)/(1 - q)$	no in general	q -exponential
Naudts exponential [10]	ϕ - ϕ -dependent	ϕ -dependent	ϕ -exponential family

Table 8. Composition laws on positive ratios and their resulting maximum-entropy distributions, on a finite outcome space with the cost J playing the role of energy. The RCL is the entry on the first row; the second row is the bilinear degree-two template before fixing $\alpha = 1$, showing how the non-degenerate bilinear coupling excludes the $\alpha = 0$ additive reduction to (C3). The squared-log surrogate is the (S1)-(S2)-respecting Cauchy-additive option in log-coordinates. Tsallis and Naudts deformations replace not only the cost but also the entropy functional, and yield non-Gibbsian distributions with power-law tails.

8.6. Scope and limitations

The finite-state setting dictates the principal limitations. Theorems 4.4–6.1 are proved on a finite outcome space Ω with strictly positive weights wherever needed; without further work they do not address continuum configuration spaces, hard-core or singular interactions, quantum reduced states, spatially extended Gibbs measures, or equivalence of ensembles in the sense of Ruelle [3] or Lanford [12]. The finite-state Lyapunov statement of Proposition 6.4 and Corollary 6.5 should not be conflated with the Boltzmann H -theorem for an interacting many-particle gas.

The framework should not be applied directly to spatially extended lattice models. For an L -site Ising or Potts system with state count $K = 2^L$ (resp. $K = q^L$), the non-asymptotic Stirling bound of Proposition 4.1 requires sample size $N \gg K \log N$, equivalently $N \gtrsim L \cdot 2^L$, far beyond what any direct enumeration on Ω can support; this is the standard curse of dimensionality for method-of-types inequalities (cf. Remark 4.2). Theorem 4.4 therefore quantitatively controls only models with K fixed, or with the active type count K^\dagger controlled along the relevant subsequence—coarse-grained or

effective-state-count models, single-molecule kinetics on a few-state graph, finite reaction networks, and the like. In coarse-grained or multiscale constructions, K should be read as the cardinality of the macrostate partition on which empirical types are accumulated; the Gibbs, entropy, and D_{KL} identities then refer to that aggregated description, and any induced Markov dynamics between macrocells must be specified separately before the relaxation statements of Section 6 apply literally. Sample-size entries such as $N_{2\sigma}$ in Section 7.6 inherit the same interpretation: they count effective independent samples at the resolution at which types are measured. Reaching spatially extended models would require either replacing Ω by a coarse-grained partition with K controlled, or replacing the multinomial counting argument of Section 4 by a transfer-matrix or variational large-deviation estimate adapted to the lattice; neither is attempted here.

8.7. Multi-constraint soft shell

The single-constraint soft-shell multinomial theorem (Theorem 4.4, equation (19)) and the exact-shell corollary for rational data (Corollary 4.5) are proved in Section 4. This subsection records the extension to several commuting affine constraints.

Several commuting linear constraints. The same soft-shell construction extends to $L \geq 1$ simultaneous affine constraints. Given real-valued cost vectors $X^{(\ell)} \in \mathbb{R}^K$ and target means $E_{\text{mc}}^{(\ell)} \in \mathbb{R}$ for $\ell = 1, \dots, L$, set

$$\begin{aligned}\Gamma &:= \left\{ p \in \Delta_K : \langle X^{(\ell)} \rangle_p = E_{\text{mc}}^{(\ell)}, \ell = 1, \dots, L \right\}, \\ \Gamma_\delta &:= \left\{ p \in \Delta_K : \max_{1 \leq \ell \leq L} |\langle X^{(\ell)} \rangle_p - E_{\text{mc}}^{(\ell)}| \leq \delta \right\},\end{aligned}$$

and let $\mathcal{F}_N^{\delta_N} := \{ \mathbf{n} \in \mathbb{Z}_{\geq 0}^K : \sum_k n_k = N, \mathbf{n}/N \in \Gamma_{\delta_N} \}$.

Proposition 8.1 (Multi-constraint soft-shell version). *Fix $L \geq 1$, $K \geq 2$, vectors $X^{(\ell)} \in \mathbb{R}^K$, and target means $E_{\text{mc}}^{(\ell)} \in \mathbb{R}$ for $\ell = 1, \dots, L$. Assume Γ is nonempty and set $h_\star^{(L)} := \max_{p \in \Gamma} H(p)$. Let $\delta_N > 0$ satisfy $\delta_N \rightarrow 0$ and $N\delta_N \rightarrow \infty$. Then $\mathcal{F}_N^{\delta_N}$ is nonempty for all sufficiently large N , and*

$$\lim_{N \rightarrow \infty} \frac{1}{N} \log \max_{\mathbf{n} \in \mathcal{F}_N^{\delta_N}} W(\mathbf{n}) = h_\star^{(L)}. \quad (28)$$

If in addition Γ has nonempty relative interior in Δ_K and $\{\mathbf{1}, X^{(1)}, \dots, X^{(L)}\}$ are linearly independent in \mathbb{R}^K , the maximizer is unique and takes the multi-parameter Gibbs form

$$p_\omega^\star = \frac{\exp\left(-\sum_\ell \beta_\ell X_\omega^{(\ell)}\right)}{Z(\boldsymbol{\beta})}, \quad Z(\boldsymbol{\beta}) := \sum_\omega \exp\left(-\sum_\ell \beta_\ell X_\omega^{(\ell)}\right), \quad (29)$$

where $\boldsymbol{\beta} = (\beta_1, \dots, \beta_L) \in \mathbb{R}^L$ is the unique solution of the moment-matching equations $\langle X^{(\ell)} \rangle_{p(\boldsymbol{\beta})} = E_{\text{mc}}^{(\ell)}$, $\ell = 1, \dots, L$.

Proof. The rate statement (28) adapts the proof of Theorem 4.4. Continuity of each $p \mapsto \langle X^{(\ell)} \rangle_p$ on the compact simplex Δ_K and continuity of H give, by the same convergent-subsequence argument as in Theorem 4.4, that $\sup_{p \in \Gamma_\delta} H(p) \rightarrow h_\star^{(L)}$ as $\delta \downarrow 0$.

For the upper bound, any $\mathbf{n} \in \mathcal{F}_N^{\delta_N}$ has empirical type $\hat{p} := \mathbf{n}/N \in \Gamma_{\delta_N}$, so Proposition 4.1 gives

$$\frac{1}{N} \log W(\mathbf{n}) \leq H(\hat{p}) + o(1) \leq \sup_{p \in \Gamma_{\delta_N}} H(p) + o(1) = h_\star^{(L)} + o(1).$$

For the lower bound, choose $p^\star \in \Gamma$ with $H(p^\star) = h_\star^{(L)}$, which exists by compactness of Γ and continuity of H . Round to $\mathbf{n}^{(N)} \in \mathbb{Z}_{\geq 0}^K$ with $\sum_k n_k^{(N)} = N$ and $\|\mathbf{n}^{(N)}/N - p^\star\|_1 \leq K/N$. Setting $M := \max_\ell \|X^{(\ell)}\|_\infty$, for each ℓ

$$|\langle X^{(\ell)} \rangle_{\mathbf{n}^{(N)}/N} - E_{\text{mc}}^{(\ell)}| = |\langle X^{(\ell)} \rangle_{\mathbf{n}^{(N)}/N} - \langle X^{(\ell)} \rangle_{p^\star}| \leq \|X^{(\ell)}\|_\infty \cdot \frac{K}{N} \leq \frac{MK}{N},$$

so $\max_\ell |\langle X^{(\ell)} \rangle_{\mathbf{n}^{(N)}/N} - E_{\text{mc}}^{(\ell)}| \leq MK/N \leq \delta_N$ for all sufficiently large N since $N\delta_N \rightarrow \infty$. Hence $\mathbf{n}^{(N)} \in \mathcal{F}_N^{\delta_N}$ eventually, and Proposition 4.1 together with continuity of H yields

$$\frac{1}{N} \log W(\mathbf{n}^{(N)}) = H(\mathbf{n}^{(N)}/N) + o(1) \longrightarrow H(p^\star) = h_\star^{(L)},$$

giving the matching liminf.

For the Gibbs form (29), H is strictly concave on Δ_K and Γ is a nonempty compact convex polytope, so the maximizer p^\star is unique. The hypothesis that Γ has nonempty relative interior in Δ_K is the standard constraint qualification ensuring existence of Karush–Kuhn–Tucker multipliers, and the linear-independence hypothesis on $\{\mathbf{1}, X^{(1)}, \dots, X^{(L)}\}$ ensures uniqueness of those multipliers (the constraint matrix has full rank $L + 1$ once the normalization $\sum_\omega p_\omega = 1$ is appended). Stationarity of the Lagrangian

$$\mathcal{L}(p, \alpha, \boldsymbol{\beta}) = - \sum_\omega p_\omega \log p_\omega - \alpha \left(\sum_\omega p_\omega - 1 \right) - \sum_\ell \beta_\ell \left(\sum_\omega p_\omega X_\omega^{(\ell)} - E_{\text{mc}}^{(\ell)} \right)$$

with respect to p_ω yields $-\log p_\omega - 1 = \alpha + \sum_\ell \beta_\ell X_\omega^{(\ell)}$, and rearranging gives $p_\omega^\star \propto \exp(-\sum_\ell \beta_\ell X_\omega^{(\ell)})$; the prefactor is fixed by normalization, producing (29). Uniqueness of $\boldsymbol{\beta}$ follows from strict convexity of the log-partition function $\boldsymbol{\beta} \mapsto \log Z(\boldsymbol{\beta})$, whose Hessian is $\text{Cov}_{p(\boldsymbol{\beta})}(X^{(1)}, \dots, X^{(L)})$; this covariance matrix is positive definite because $\{X^{(\ell)} - \langle X^{(\ell)} \rangle_{p(\boldsymbol{\beta})} \mathbf{1}\}_\ell$ are linearly independent in \mathbb{R}^K under the hypothesis on $\{\mathbf{1}, X^{(1)}, \dots, X^{(L)}\}$. \square

8.8. Numerical illustration with irrational costs

The engineered ratios of Section 7 were chosen so that the RCL costs $J(r_\omega)$ are rational and Corollary 4.5 applies directly. To see Theorem 4.4 in operation in the generic case, replace those ratios by

$$r_2 = e^{-1}, \quad r_3 = e^{-3/2},$$

so the cost vector becomes

$$X = (0, \cosh 1 - 1, \cosh(3/2) - 1) \approx (0, 0.5431, 1.3524),$$

with both nonzero entries irrational. The uniform-cost reference is $\langle X \rangle_{(1/3, 1/3, 1/3)} = \frac{1}{3} [\cosh 1 + \cosh(3/2)] - 1 \approx 0.6318$, so the admissible mean-cost range under positive temperature is $E_{\text{mc}} \in (0, 0.6318)$. Fix $E_{\text{mc}} = 0.50$, an interior value comparable to the rational example. The moment-matching equation $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$ is solved by $\beta^* \approx 0.4462$, $T_{\text{R}} = 1/\beta^* \approx 2.241$, and yields the analytic Gibbs measure

$$p^* \approx (0.4289, 0.3366, 0.2346), \quad Z(\beta^*) \approx 2.3317, \quad H(p^*) \approx 1.0697.$$

The constraint $\sum_k n_k X_k = N E_{\text{mc}}$ admits no integer solution for any N because the entries of X are irrationally related, so the exact shell \mathcal{F}_N is always empty. Theorem 4.4 still applies: take the tolerance $\delta_N = N^{-1/2}$, which satisfies $\delta_N \rightarrow 0$ and $N\delta_N \rightarrow \infty$, and round Np^* to the nearest integer composition. At $N = 10^3$ this gives the type $\mathbf{n}^{(N)} = (429, 337, 234)$ with

$$|\langle X \rangle_{\mathbf{n}^{(N)}/N} - E_{\text{mc}}| \approx 5.18 \times 10^{-4} \ll \delta_{10^3} \approx 3.16 \times 10^{-2},$$

well within the bound $K\|X\|_{\infty}/N \approx 4.06 \times 10^{-3}$ from the proof of the proposition. The empirical entropy $H(\mathbf{n}^{(N)}/N) \approx 1.0695$ differs from $h_{\star} = H(p^*)$ by 2.32×10^{-4} , consistent with the $N^{-1} \log W(\mathbf{n}) \rightarrow h_{\star}$ statement (19) at this sample size. The same construction supplies a feasible soft-shell type for every $N \geq \lceil K\|X\|_{\infty}/\delta_N \rceil = \lceil 4.06 N^{1/2} \rceil$, i.e. for all $N \geq 17$ at this tolerance schedule. Multi-constraint examples are handled identically by replacing X by the constraint matrix $(X^{(\ell)})_{\ell=1}^L$ and applying Proposition 8.1.

9. Conclusion

Under axioms (A1)–(A3), the cost J in (5) is determined (Theorem 1.1), multinomial counting yields Shannon entropy as the rate functional for empirical types (Theorem 4.4 and Corollary 4.5, with explicit non-asymptotic Stirling bound, Proposition 4.1, and the multi-constraint extension, Proposition 8.1), and convex maximization at fixed mean cost gives the finite-state Gibbs weights (Proposition 5.3) together with the free-energy–Kullback–Leibler identity (Proposition 6.1). Lemma 2.1 and Proposition 3.1 supply the companion ratio-level and divergence-level bounds; Proposition 6.4 and Corollary 6.5 extend the static identities to detailed-balance relaxation.

The contribution is a finite-state, explicitly delimited construction linking a ratio-cost axiom to the classical canonical-ensemble machinery through the thread $\{r_{\omega}\} \rightarrow J \rightarrow X \rightarrow (\text{Gibbs}, F_{\text{R}}, D_{\text{KL}})$; Steps 4–9 of Table 7 are standard, and the RCL itself remains a modeling axiom rather than a consequence of detailed balance alone. The Tsallis comparison of Section 7.5 and the alternative-composition-law tabulation of Section 8.5 show that the d’Alembert law (2) is one choice among several, motivated by the bookkeeping symmetries (S1)–(S3) but not forced by them.

Natural continuations include spatially extended models and equivalence of ensembles, continuum configuration spaces with singular interactions, dynamical analogs of the J -cost bounds beyond the finite-state Lyapunov statements, and tighter control of the curse-of-dimensionality regime where the state count K grows with system size (cf. Remark 4.2). The axiomatic question raised in Section 8.2—whether the Shore–Johnson axioms restated in ratio coordinates force, or merely permit, the d’Alembert composition law (2)—is also left open here.

The methodological contribution is the organization of these ingredients into a self-contained route from dimensionless ratios to a canonical exponential family on a finite outcome space. The cost-classification step (N1) follows Theorem 1.1 under measurability and local integrability as in Washburn–Zlatanović [13]; the Stirling bound (N2) is a direct application of Robbins’ 1955 inequalities to the finite-state method of types; Theorem 4.4 and Corollary 4.5 organize the soft-shell and exact-shell counting statements; and the unified notation (N3), together with the alternative-composition tabulation and Tsallis comparison, makes explicit how the chain $\{r_\omega\} \rightarrow J \rightarrow X \rightarrow (\text{Gibbs}, F_R, D_{\text{KL}})$ depends on the chosen ratio-cost postulates. In this scope, the paper supplies a finite-state construction for systems whose primitives are dimensionless ratios, with approximation errors quantified and alternative composition laws tabulated. The sample-size power calculation of Section 7.6 quantifies how distinguishable the RCL prediction is from the natural “affinity-as-energy” and Tsallis alternatives at fixed mean cost, conditional on the RCL Gibbs law as generating distribution; comparison against measured data would require an experimental three-state system with independently measured rate ratios and stationary populations. The construction does not attempt to derive the axioms (S1)–(S3) or the analytic-degree-2 closure from a more primitive principle: it converts those modeling commitments, together with the input ratios, into a closed-form Gibbs-like measure and the associated free-energy / Kullback–Leibler structure on a finite outcome space.

Author contributions

J.W.: conceptualization of the project, original draft of the manuscript, analysis, and methodology (CRediT: Conceptualization, Formal Analysis, Investigation, Methodology, Writing – Original Draft). M.S.: formal analysis, manuscript revision, validation of derivations and numerical examples, project administration, and writing–review and editing (CRediT: Formal Analysis, Validation, Project Administration, Writing – Review & Editing).

Funding

This work received no external financial support.

Conflict of interest

The authors declare that they have no known competing financial interests or personal relationships that could have appeared to influence the work reported in this paper.

Data and code availability

No experimental data were generated or analyzed. Short Python/NumPy/SciPy scripts that reproduce the numerical entries in Tables 2–5 and Figures 1–2 accompany this manuscript as separate files together with the `.tex` source: they solve the moment-matching equation $\langle X \rangle_{p^{(\beta)}} = E_{\text{mc}}$ numerically and evaluate the Gibbs weights, Tsallis q_T -exponential weights, free energy, divergences, and sample-size power-calculation predictions in closed form. The scripts are included in the arXiv ancillary source bundle.

Acknowledgments

The authors thank colleagues at the Recognition Physics Institute for helpful discussions during the development of this work.

References

- [1] Ludwig Boltzmann. Über die Beziehung zwischen dem zweiten Hauptsatze der mechanischen Wärmetheorie und der Wahrscheinlichkeitsrechnung respektive den Sätzen über das Wärmegleichgewicht. *Wiener Berichte*, 76:373–435, 1877.
- [2] J. Willard Gibbs. *Elementary Principles in Statistical Mechanics*. Yale University Press, 1902.
- [3] David Ruelle. *Statistical Mechanics: Rigorous Results*. W. A. Benjamin, 1969.
- [4] E. T. Jaynes. Information theory and statistical mechanics. *Phys. Rev.*, 106:620–630, 1957.
- [5] E. T. Jaynes. Information theory and statistical mechanics. II. *Phys. Rev.*, 108:171–190, 1957.
- [6] Steve Presse, Kingshuk Ghosh, Julian Lee, and Ken A. Dill. Principles of maximum entropy and maximum caliber in statistical physics. *Rev. Mod. Phys.*, 85:1115–1141, 2013.
- [7] John E. Shore and Rodney W. Johnson. Axiomatic derivation of the principle of maximum entropy and the principle of minimum cross-entropy. *IEEE Trans. Inform. Theory*, 26(1):26–37, 1980.
- [8] Ariel Caticha. *Entropic Inference and the Foundations of Physics*. EBL-Schweitzer (USP / EBEB), 2012. Lecture notes available at <https://www.albany.edu/physics/faculty/ariel-caticha>.
- [9] Constantino Tsallis. Possible generalization of Boltzmann–Gibbs statistics. *J. Stat. Phys.*, 52(1–2):479–487, 1988.
- [10] Jan Naudts. *Generalised Thermostatistics*. Springer, 2011.
- [11] Hugo Touchette. The large deviation approach to statistical mechanics. *Phys. Rep.*, 478:1–69, 2009.
- [12] Oscar E. Lanford. Entropy and equilibrium states in classical statistical mechanics. In *Statistical Mechanics and Mathematical Problems*, volume 20 of *Lecture Notes in Physics*, pages 1–113. Springer, 1973.
- [13] Jonathan Washburn and Milan Zlatanović. Uniqueness of the canonical reciprocal cost. *Mathematics*, 14(6):935, 2026. Also available as arXiv:2602.05753 [math.CA].
- [14] J. Aczél. *Lectures on Functional Equations and Their Applications*. Academic Press, 1966.
- [15] Marek Kuczma, Bogdan Choczewski, and Roman Ger. *Iterative Functional Equations*, volume 32 of *Encyclopedia of Mathematics and Its Applications*. Cambridge University Press, 2009. Reprint of the 1990 original.
- [16] Jonathan Washburn and Amir Rahnamai Barghi. The coercive projection theorem for canonical reciprocal costs, 2026. Preprint, arXiv:2603.20205.
- [17] Herbert Robbins. A remark on Stirling’s formula. *Amer. Math. Monthly*, 62(1):26–29, 1955.
- [18] Richard S. Ellis. *Entropy, Large Deviations, and Statistical Mechanics*. Springer, 1985.
- [19] J. Schnakenberg. Network theory of microscopic and macroscopic behavior of master equation systems. *Rev. Mod. Phys.*, 48:571–585, 1976.
- [20] Thomas M. Cover and Joy A. Thomas. *Elements of Information Theory*. Wiley, 2 edition, 2006.
- [21] Imre Csiszár and János Körner. *Information Theory: Coding Theorems for Discrete Memoryless Systems*. Academic Press, 1981.
- [22] Amir Dembo and Ofer Zeitouni. *Large Deviations Techniques and Applications*. Springer, 2 edition, 2010. Corrected reprint of the 1998 edition.
- [23] David A. Levin, Yuval Peres, and Elizabeth L. Wilmer. *Markov Chains and Mixing Times*. American Mathematical Society, 2 edition, 2017.
- [24] Dario Lucente, Marco Baldovin, Andrea Puglisi, and Angelo Vulpiani. H-theorem at negative temperature: the random exchange model with bounds. *J. Stat. Mech.*, 2025(1):013210, 2025.
- [25] Solomon Kullback and Richard A. Leibler. On information and sufficiency. *Ann. Math. Statist.*, 22:79–86, 1951.
- [26] Imre Csiszár. Information-type measures of difference of probability distributions and indirect observations. *Studia Sci. Math. Hungar.*, 2:299–318, 1967.
- [27] S. M. Ali and S. D. Silvey. A general class of coefficients of divergence of one distribution from another. *J. Roy. Statist. Soc. Ser. B Methodol.*, 28(1):131–142, 1966.
- [28] Alexander Schrijver. *Theory of Linear and Integer Programming*. Wiley, 1998.

- [29] I. N. Sanov. On the probability of large deviations of random magnitudes. *Mat. Sb. (N.S.)*, 42(1):11–44, 1957.
- [30] Jürgen Voigt. Stochastic operators, information, and entropy. *Comm. Math. Phys.*, 81:31–38, 1981.
- [31] J. R. Norris. *Markov Chains*. Cambridge Series in Statistical and Probabilistic Mathematics. Cambridge University Press, 1997.
- [32] Udo Seifert. Stochastic thermodynamics, fluctuation theorems and molecular machines. *Rep. Prog. Phys.*, 75(12):126001, 2012.
- [33] Riccardo Rao and Massimiliano Esposito. Nonequilibrium thermodynamics of chemical reaction networks: Wisdom from stochastic thermodynamics. *Phys. Rev. X*, 6:041064, 2016.
- [34] Ariel Caticha. Relative entropy and inductive inference. In *Bayesian Inference and Maximum Entropy Methods in Science and Engineering*, volume 707 of *AIP Conference Proceedings*, pages 75–96. American Institute of Physics, 2004. Also available as arXiv:physics/0311093.
- [35] Shun-ichi Amari. *Differential-Geometrical Methods in Statistics*, volume 28 of *Lecture Notes in Statistics*. Springer, 1985.
- [36] Shun-ichi Amari and Hiroshi Nagaoka. *Methods of Information Geometry*, volume 191 of *Translations of Mathematical Monographs*. American Mathematical Society, 2000.
- [37] Fumitada Itakura and Shuzo Saito. Analysis synthesis telephony based on the maximum likelihood method. *Reports of the 6th International Congress on Acoustics*, pages C–17–C–20, 1968.
- [38] Zeineb Chebbi and Maher Moakher. Means of Hermitian positive-definite matrices based on the log-determinant α -divergence function. *Linear Algebra Appl.*, 436:1872–1889, 2012.
- [39] Suvrit Sra. A new metric on the manifold of kernel matrices with application to matrix geometric means. *Adv. Neural Inf. Process. Syst.*, 25:144–152, 2012.
- [40] János Aczél and Jean Dhombres. *Functional Equations in Several Variables*, volume 31 of *Encyclopedia of Mathematics and its Applications*. Cambridge University Press, 1989.
- [41] Massimiliano Esposito and Christian Van den Broeck. Three faces of the second law. I. Master equation formulation. *Phys. Rev. E*, 82:011143, 2010.
- [42] Matteo Polettini and Massimiliano Esposito. Irreversible thermodynamics of open chemical networks. I. Emergent cycles and broken conservation laws. *J. Chem. Phys.*, 141(2):024117, 2014.
- [43] Sosuke Ito. Stochastic thermodynamic interpretation of information geometry. *Phys. Rev. Lett.*, 121:030605, 2018.
- [44] Andre C. Barato and Udo Seifert. Thermodynamic uncertainty relation for biomolecular processes. *Phys. Rev. Lett.*, 114:158101, 2015.
- [45] Jan Naudts. Estimators, escort probabilities, and ϕ -exponential families in statistical physics. *J. Inequal. Pure Appl. Math.*, 5(4):Article 102, 2004.