The Many Faces of the Partition Function

A mini-series for graduate students crossing between mathematics and theoretical physics

Abstract

These notes follow a single object — the partition function — through five re-tellings: the statistical-mechanical *sum over states*, the quantum-statistical *trace*, the field-theoretic *path integral*, the conformal *torus amplitude*, and the representation-theoretic *character*. A short interlude abstracts the Boltzmann factor that underlies all of them, and a closing chapter reframes the whole story categorically. Two concrete examples — the free boson and the free fermion — run through every chapter and converge, at the end, on superconformal field theory. The aim is translation: to let a reader fluent in one of these dialects hear the others as the same sentence.

Contents

Orientation: one idea in five disguises

Ask a condensed-matter physicist, a field theorist, and a representation theorist what a *partition function* is, and you will get three answers that sound unrelated. The first will write a sum of Boltzmann weights; the second a path integral; the third a generating function counting dimensions of graded pieces of a module. These notes are an argument that they are describing the same object from different angles, and that learning to translate between the descriptions is most of what it means to move comfortably between the fields.

The one idea

> *The partition function is a weighted trace, which is a graded character. Its deep payoff is modular invariance, which binds the physical torus to the mathematical characters.*

We will meet the same object wearing five costumes, one per chapter:

$$\underbrace{Z=\sum_i \mathrm{e}^{-\beta E_i}}_{\text{L1: a sum}} \longrightarrow \underbrace{Z=\mathrm{Tr}\,\mathrm{e}^{-\beta H}}_{\text{L2: a trace}} \longrightarrow \underbrace{Z=\int \mathcal{D}\phi\,\mathrm{e}^{-S_E}}_{\text{L3: a path integral}} \longrightarrow \underbrace{Z(\tau)=\mathrm{Tr}\,q^{L_0-\frac{c}{24}}\,\bar{q}^{\bar{L}_0-\frac{c}{24}}}_{\text{L4: a torus amplitude}} \longrightarrow \underbrace{\mathrm{ch}\,V=\sum_n (\dim V_n)\,q^n}_{\text{L5: a character}}$$

Three threads to watch

Thread 1 — the Boltzmann factor, abstracted

The weight $\mathrm{e}^{-\beta E}$ is the most-used and least-examined object in the subject. Its abstraction — *a homomorphism turning additive conserved charges into multiplicative weights, with one conjugate potential per charge* — is developed in Lectures 1–2 and consolidated in the Interlude. Watch the exponent grow new terms (a chemical potential, a sign grading) and watch those conjugate potentials get renamed downstream: $\beta \to \mathrm{Im}\,\tau$, fugacity $\to$ elliptic variable y, $\to$ the Cartan variables of a character.

Thread 2 — the boson and the fermion

Introduced as single modes in Lecture 1, they are promoted at each stage until, on the torus, the boson yields the Dedekind eta and the fermion's four spin structures yield the Jacobi theta functions. That is precisely the data needed to reach superconformal field theory — NS/R sectors, the Witten index, the elliptic genus — at the close.

Thread 3 — the Casimir constant $-c/24$

A single quantity surfaces three times: as the regularized vacuum energy $\zeta(-1)=-\frac{1}{12}$ in Lecture 3, as the $q^{1/24}$ in $\eta(\tau)$ in Lecture 4, and as the universal $L_0-\frac{c}{24}$ grading shift in Lecture 5.

Conventions, fixed once

- We work in **Euclidean signature** for partition functions; Lorentzian language appears only as motivation.
- **The symbol q changes meaning, and we will say so each time.** In thermal physics a single mode contributes with $q = \mathrm{e}^{-\beta\hbar\omega}$; on the torus $q = \mathrm{e}^{2\pi \mathrm{i}\tau}$. The two agree only after β is identified with an imaginary period.
- $\beta = 1/(k_B T)$; we set $k_B = \hbar = 1$ after Lecture 2.
- $\tau = \tau_1 + \mathrm{i}\tau_2 \in \mathbb{H}$, the upper half-plane, so $\operatorname{Im}\tau = \tau_2 > 0$.
- In the 2d chapters we take $c = \bar{c}$: one central charge for both chiralities.
- Vocabulary clashes are flagged inline: “weight” (CFT conformal weight h vs. Lie-theory weight λ), “character” (group trace vs. graded dimension), “level” (affine k vs. an energy level).

Each chapter opens with a sidebar—“For the reader from camp X”—mapping the vocabulary you already own onto the one we are about to use. Heavy machinery and proofs are quarantined in “For the curious” boxes; nothing in the main narrative depends on them.

Chapter 1

Statistical Mechanics: the partition function as a sum over states

For the reader from mathematics

Read β as a formal variable and $Z(\beta) = \sum_i \mathrm{e}^{-\beta E_i}$ as a generating function: a Laplace/Dirichlet-type transform of the energy spectrum. Every thermodynamic quantity is a logarithmic derivative of it. The only "physics" input is the principle that selects the weights $\mathrm{e}^{-\beta E_i}$, which we derive rather than postulate.

The partition function is born as a normalizing constant. What makes it the central object of the theory is that, once you have it, you have everything: the entire thermodynamics of a system is encoded in the single function $\log Z$.

1.1 Where the Boltzmann weight comes from

Fix a system with microstates i of energy E_i. We do not know which microstate the system occupies; we know only its *mean* energy. Among all probability distributions $\{p_i\}$ with $\sum_i p_i = 1$ and $\sum_i p_i E_i = \langle E \rangle$ fixed, which should we choose? The principle of maximum entropy says: the least-committal one, maximizing the Gibbs–Shannon entropy $S = -\sum_i p_i \log p_i$. Lagrange multipliers α, β enforce the two constraints; we extremize

$$-\sum_i p_i \log p_i - \alpha\Big(\sum_i p_i - 1\Big) - \beta\Big(\sum_i p_i E_i - \langle E \rangle\Big) \tag{1.1}$$

over each p_i in turn. This gives one first-order condition per state i—an ordinary unconstrained calculus problem, one equation per unknown—namely $\log p_i = -1 - \alpha - \beta E_i$, that is

$$\boxed{p_i = \frac{\mathrm{e}^{-\beta E_i}}{Z}, \qquad Z = \sum_i \mathrm{e}^{-\beta E_i}.} \tag{1.2}$$

The multiplier β enforcing the energy constraint is the inverse temperature, $\beta = 1/(k_B T)$, as we will see shortly. The normalizer Z is the **partition function**: it is what *partitions* unit probability among the states. Nothing was postulated about exponentials; the exponential is forced by maximizing entropy at fixed mean energy. We return to this point, hard, in the Interlude.

1.2 Everything is a derivative of $\log Z$

Because Z packages the whole spectrum, its logarithm generates the thermodynamics:

$$\langle E\rangle = \frac{\sum_i E_i \mathrm{e}^{-\beta E_i}}{Z} = -\frac{\partial \log Z}{\partial \beta}, \tag{1.3}$$

$$\mathrm{Var}(E) = \langle E^2\rangle - \langle E\rangle^2 = \frac{\partial^2 \log Z}{\partial \beta^2}, \tag{1.4}$$

so $\log Z$ is literally the cumulant generating function of the energy.

Where does the identification $\beta = 1/(k_B T)$ from §1.1 actually come from? Evaluate the informational entropy on the maximizing distribution itself: since $\log p_i = -\beta E_i - \log Z$,

$$S = -\sum_i p_i \log p_i = \beta\langle E\rangle + \log Z. \tag{1.5}$$

Multiply by k_B to give S thermodynamic units and differentiate with respect to $\langle E\rangle$ at fixed volume and particle number, using $\partial \log Z/\partial\beta = -\langle E\rangle$ from above—the $\partial\beta/\partial\langle E\rangle$ terms cancel exactly, leaving

$$k_B \frac{\partial S}{\partial\langle E\rangle} = k_B \beta. \tag{1.6}$$

Classical thermodynamics *defines* temperature by $\partial S/\partial E = 1/T$; comparing the two fixes $\beta = 1/(k_B T)$ exactly. The Lagrange multiplier of §1.1 and the thermodynamic inverse temperature are the same object because both are, at bottom, $\partial S/\partial E$.

The free energy is $F = -\beta^{-1}\log Z$, the entropy $S = -\partial F/\partial T$, the heat capacity $C = \partial\langle E\rangle/\partial T$, and so on. One function, differentiated, yields the lot.

For the reader from physics

If you already know all this, the only new emphasis is structural: keep your eye on (1.2) as a *map* from a spectrum to a number, and on the multiplicativity of §1.4. Those two features—and not the thermodynamics—are what survive into quantum statistics, field theory, and representation theory.

1.3 Two elementary systems: the boson and the fermion appear

Example 1.1 (Two-level system)**.** A single degree of freedom with energies $\{0, \epsilon\}$ has $Z = 1+\mathrm{e}^{-\beta\epsilon}$. Its heat capacity peaks at $k_B T \sim \epsilon$ (the Schottky anomaly)[1]: the system can only "notice" temperature once $k_B T$ is comparable to its gap.

Example 1.2 (A single harmonic oscillator — the *boson*)**.** With $E_n = \hbar\omega(n+\frac{1}{2})$, $n = 0, 1, 2, \ldots$, the geometric series sums to

$$Z_\mathrm{b} = \sum_{n\ge 0} \mathrm{e}^{-\beta\hbar\omega(n+1/2)} = \frac{\mathrm{e}^{-\beta\hbar\omega/2}}{1-\mathrm{e}^{-\beta\hbar\omega}} = \frac{1}{2\sinh(\beta\hbar\omega/2)}. \tag{1.7}$$

A bosonic mode admits *any* occupation number $n \ge 0$; that is why we get a full geometric series.[2]

[1] Explicitly, $C(T) = k_B\left(\frac{\epsilon}{k_B T}\right)^2 \frac{\mathrm{e}^{-\epsilon/k_B T}}{(1+\mathrm{e}^{-\epsilon/k_B T})^2}$—a bump vanishing both as $T \to 0$ (frozen out) and $T \to \infty$ (equal population, no distinguishing power), unlike the monotonic heat capacity of a classical system. Named for Walter Schottky (1922); such bumps have since been used to diagnose gapped low-lying levels—impurity spins, nuclear hyperfine levels, tunneling two-level systems in glasses—directly from bulk specific-heat data.

[2] Readers who know the Feynman path integral will recognize (1.7) lurking inside the harmonic-oscillator propagator, $K(x_b, t; x_a, 0) \propto (\sin\omega t)^{-1/2}\exp(\mathrm{i}\cdots)$: continue to imaginary time, $t \to -\mathrm{i}\beta\hbar$ (so $\sin\omega t \to -\mathrm{i}\sinh\beta\hbar\omega$), set $x_b = x_a$, and integrate over x_a—that is, trace over *closed* paths—and the Gaussian integral yields exactly $1/(2\sinh(\beta\hbar\omega/2))$. No accident: $\mathrm{e}^{-\beta H}$ is the time-evolution operator at imaginary time, so the thermal Z is the Feynman amplitude to return, summed over starting points. Lecture 3 is this footnote, industrialized.

The $-c/24$ thread starts here

The factor $\mathrm{e}^{-\beta\hbar\omega/2}$ in (1.7) is the **zero-point energy** $\frac{1}{2}\hbar\omega$ of the oscillator. It looks like a harmless prefactor now. Summed over infinitely many modes it will diverge and get regularized in L3 via $\sum n \overset{\zeta}{=} -\frac{1}{12}$ (the energy itself carrying the extra $\frac{1}{2}$: $-\frac{1}{24}$ per unit of mode spacing), become the $q^{1/24}$ of the Dedekind eta in L4, and end as the $-c/24$ grading shift in L5.

Example 1.3 (A single fermionic mode — the *fermion*)**.** With occupation $n \in \{0,1\}$ (Pauli exclusion) and energy $E = \epsilon n$,

$$Z_{\mathrm{f}} = 1 + \mathrm{e}^{-\beta\epsilon}. \tag{1.8}$$

The contrast with (1.7) is the whole of quantum statistics in miniature: bosons sum a geometric series over all n, fermions truncate at two terms.

1.4 Independence becomes multiplication

If a system splits into non-interacting parts with additive energy $E_{(i,j)} = E_i^A + E_j^B$, then

$$Z = \sum_{i,j} \mathrm{e}^{-\beta(E_i^A+E_j^B)} = \Big(\sum_i \mathrm{e}^{-\beta E_i^A}\Big)\Big(\sum_j \mathrm{e}^{-\beta E_j^B}\Big) = Z_A\, Z_B, \tag{1.9}$$

and more generally $Z = \prod_k Z_k$, $\log Z = \sum_k \log Z_k$. This innocuous identity is the seed of *every* product formula in these notes.

The Boltzmann factor as a homomorphism

Look closely at (1.9). The single fact that powers it is

$$\mathrm{e}^{-\beta(E_1+E_2)} = \mathrm{e}^{-\beta E_1}\,\mathrm{e}^{-\beta E_2}. \tag{1.10}$$

The map $E \mapsto \mathrm{e}^{-\beta E}$ sends *addition* of energies to *multiplication* of weights. That is the defining property of a one-dimensional *character* of the additive group of energies—and it is why independent systems multiply, why Z factorizes over modes, and (eventually) why η and θ are infinite products. We will also soon want to weight states by a second additive quantity, a conserved charge N, via $\mathrm{e}^{-\beta(E-\mu N)}$: the new knob μ is the **chemical potential**, developed in Lecture 2.

For the curious: the symplectic route to $\beta = 1/(k_B T)$

The identification $\beta = 1/(k_B T)$ above leaned on classical thermodynamics as an external input. There is a purely geometric route that never invokes it. Phase space carries a Liouville volume, preserved by Hamiltonian flow; let $\Omega(E) = \mathrm{Vol}\{H \le E\}$ be the enclosed volume below energy E—a symplectic quantity, nothing thermal about it yet. Define the microcanonical entropy $S(E) = \log\Omega(E)$ and *define* $\beta := dS/dE$: the logarithmic growth rate of a volume, pure calculus so far. This β agrees with the Lagrange multiplier of §1.1: for large systems $Z(\beta) = \int dE\,\Omega'(E)\,\mathrm{e}^{-\beta E}$ is dominated by the energy E^* where $\log\Omega'(E) - \beta E$ is stationary, i.e. $\beta = d\log\Omega'/dE|_{E^*}$; and since the derivatives of $\log\Omega$ and $\log\Omega'$ agree up to subextensive corrections, Laplace's method makes the two definitions coincide in the thermodynamic limit. The one physical, not purely geometric, ingredient left is the postulate that equilibrium maximizes accessible phase volume: two weakly coupled systems distribute their energy so as to maximize $\Omega'_A(E_A)\,\Omega'_B(E_{\mathrm{tot}} - E_A)$, which forces equal β at the maximum—the statement that underwrites the zeroth law, with lighter machinery than the full Clausius–Kelvin apparatus. The same Liouville volume returns as the path-integral measure in Lecture 3.

Rosetta notes

- $\beta \leftrightarrow$ inverse temperature, soon to be reinterpreted as an imaginary time and then as $\operatorname{Im}\tau$.
- $\log Z \leftrightarrow$ the free energy, the generating function of energy cumulants.
- $Z = \prod_k Z_k \leftrightarrow$ the product formulas $\prod(1-q^n)^{-1}$ and $\prod(1+q^n)$ that become η and θ.
- The sum $\sum_i$ over a labelled basis is about to become a basis-independent trace.

Chapter 2

Quantum Statistics: the partition function as a trace

For the reader from mathematics

The upgrade here is from a sum over a labelled index set to a *trace* on $\mathrm{End}(\mathcal{H})$. A trace is basis-independent and cyclic; those two properties are doing real work. The Fock-space examples turn the geometric and finite series of Lecture 1 into the infinite products $\prod(1-q^n)^{-1}$ and $\prod(1+q^n)$ — the same series a number theorist meets as generating functions for integer partitions.

2.1 From a sum to a trace

Quantum mechanically the energies E_i are eigenvalues of a Hamiltonian operator H on a Hilbert space $\mathcal{H}$. The Boltzmann factor is no longer a number attached to a configuration; it is the *operator* $\mathrm{e}^{-\beta H}$, and the partition function is its trace:

$$\boxed{Z = \mathrm{Tr}_{\mathcal{H}}\, \mathrm{e}^{-\beta H} = \sum_i \mathrm{e}^{-\beta E_i}.} \tag{2.1}$$

Evaluating the trace in the energy eigenbasis recovers Lecture 1. But the trace is *basis-independent*: we may compute it in any basis, and that freedom is exactly what makes Z a robust object rather than an artifact of a chosen labelling.

The L1 → L2 distinction — worth dwelling on

Students routinely conflate "statistical mechanics" and "quantum statistics." The difference is precisely the promotion of the Boltzmann factor from a *scalar* $\mathrm{e}^{-\beta E_i}$ (L1) to an *operator* $\mathrm{e}^{-\beta H}$ (L2). Three consequences follow, none cosmetic:

1. With non-commuting observables there is no preferred basis, so *basis-independence* — the trace — is what makes Z well-defined.
2. Indistinguishability of identical quanta forces the Bose vs. Fermi alternatives (§2.2).
3. $\mathrm{e}^{-\beta H}$ is imaginary-time evolution, $\mathrm{e}^{-\beta H} = \mathrm{e}^{-\mathrm{i}Ht}\big|_{t=-\mathrm{i}\beta}$, the hinge that swings us into the path integral of Lecture 3.

The (unnormalized) state of the system is the density operator $\rho = \mathrm{e}^{-\beta H}/Z$, and expectation values are $\langle\mathcal{O}\rangle = \mathrm{Tr}(\rho\,\mathcal{O})$. Cyclicity of the trace, $\mathrm{Tr}(AB) = \mathrm{Tr}(BA)$, is the algebraic seed of the Kubo–Martin–Schwinger condition we will meet in L3 and of modular invariance in L4.

2.2 Fock space and the two universal products

The trace (2.1) may be computed in *any* basis; for a collection of non-interacting quantum modes the convenient one is the **occupation-number** basis of a **Fock space**. We build it one mode at a time, and in that basis the trace is a sum of diagonal matrix elements $\langle n|\mathrm{e}^{-\beta H}|n\rangle$—each basis ket sandwiched around the Boltzmann operator.

A single bosonic mode. Write the mode with ladder operators $[a, a^\dagger] = 1$. The number operator $\hat{n} = a^\dagger a$ has the tower of eigenstates $|n\rangle$, $n = 0, 1, 2, \ldots$, generated from the vacuum by $a^\dagger|n\rangle = \sqrt{n+1}\,|n+1\rangle$ (and lowered by $a|n\rangle = \sqrt{n}\,|n-1\rangle$). With $H = \hbar\omega(\hat{n} + \frac{1}{2})$ these are energy eigenstates, so each sandwich $\langle n|\mathrm{e}^{-\beta H}|n\rangle$ returns the eigenvalue $\mathrm{e}^{-\beta E_n}$ and the trace is a bare sum of Boltzmann numbers:

$$\mathrm{Tr}_{\mathcal{H}}\ \mathrm{e}^{-\beta H} = \sum_{n\geq 0}\langle n|\mathrm{e}^{-\beta H}|n\rangle = \sum_{n\geq 0}\mathrm{e}^{-\beta\hbar\omega(n+1/2)} = \frac{\mathrm{e}^{-\beta\hbar\omega/2}}{1-\mathrm{e}^{-\beta\hbar\omega}}. \tag{2.2}$$

This is (1.7) of Lecture 1, now read as a trace. A bosonic mode admits any occupation $n \geq 0$: the tower is infinite.

A single fermionic mode. Now the operators *anticommute*, $\{c, c^\dagger\} = 1$ and $c^2 = 0$. That $c^2 = 0$ is the Pauli principle, and it truncates the tower to just two states, $|0\rangle$ and $|1\rangle = c^\dagger|0\rangle$: the Fock space is two-dimensional. With $H = \epsilon\, c^\dagger c$ the trace is a two-term sum,

$$\mathrm{Tr}_{\mathcal{H}}\ \mathrm{e}^{-\beta H} = \langle 0|\mathrm{e}^{-\beta H}|0\rangle + \langle 1|\mathrm{e}^{-\beta H}|1\rangle = 1 + \mathrm{e}^{-\beta\epsilon}, \tag{2.3}$$

recovering (1.8). Bose versus Fermi is nothing but the length of the tower the trace runs over—an infinite geometric series against two terms.

For the reader from mathematics

Fock space is built from the single-particle space $V = \bigoplus_k \mathbb{C}$ (one dimension per mode) by a functor: bosonic Fock space is the symmetric algebra $\mathrm{Sym}(V) = \bigoplus_n \mathrm{Sym}^n V$, fermionic Fock space is the exterior algebra $\bigwedge V = \bigoplus_n \bigwedge^n V$, and occupation number is exactly the symmetric/exterior grading. This is not a passing analogy: in Lecture 5 $\mathrm{Sym}(V)$ becomes the Heisenberg vertex algebra and $\bigwedge V$ a Clifford (super) vertex algebra, generated by exactly the modes assembled here.

Many modes: the trace factorizes. Assemble all the modes. The Hilbert space is the tensor product $\mathcal{H} = \bigotimes_k \mathcal{H}_k$, with occupation-number basis $|\{n_k\}\rangle = \bigotimes_k |n_k\rangle$; the Hamiltonian is additive, $H = \sum_k H_k$; and because the H_k act on separate factors, $\mathrm{e}^{-\beta H} = \bigotimes_k \mathrm{e}^{-\beta H_k}$. A diagonal matrix element therefore splits into a product over modes, $\langle\{n_k\}|\mathrm{e}^{-\beta H}|\{n_k\}\rangle = \prod_k \langle n_k|\mathrm{e}^{-\beta H_k}|n_k\rangle$, and summing over every configuration $\{n_k\}$ turns a sum of products into a product of sums:

$$\begin{aligned} Z = \mathrm{Tr}\ \mathrm{e}^{-\beta H} &= \sum_{\{n_k\}}\prod_k \langle n_k|\mathrm{e}^{-\beta H_k}|n_k\rangle \\ &= \prod_k\Big(\sum_{n_k}\langle n_k|\mathrm{e}^{-\beta H_k}|n_k\rangle\Big) = \prod_k Z_k. \end{aligned} \tag{2.4}$$

Exchanging "sum over all configurations" for "product over modes" is the operator version of the factorization (1.9): independent modes multiply, and each factor Z_k is a single-mode trace of the kind just computed.

Now specialize to a tower of modes with energies $\hbar\omega_n = n\,\hbar\omega$, $n \geq 1$, and write $q = \mathrm{e}^{-\beta\hbar\omega}$. **(Note: this is the thermal q, not yet the torus q.)** Mode n then carries the weight q^n

per quantum. Measuring energies from each mode's ground state, its bosonic factor (2.2) is $\sum_{m\geq 0} q^{nm} = (1-q^n)^{-1}$ and its fermionic factor (2.3) is $1+q^n$. The product (2.4) over the whole tower is then

$$\text{Bosons (each } n \text{ occupied } 0,1,2,\ldots \text{ times):} \quad Z_{\mathrm{b}} = \prod_{n\geq 1} \frac{1}{1-q^n}, \tag{2.5}$$

$$\text{Fermions (each } n \text{ occupied } 0 \text{ or } 1 \text{ time):} \quad Z_{\mathrm{f}} = \prod_{n\geq 1} (1+q^n). \tag{2.6}$$

These two products—one for each statistics—are the protagonists of the rest of the series.

The boson/fermion thread, now infinite products

Equation (2.5) is, up to the factor $q^{-1/24}$ (the zero-point energies of Example 1.2 summed over the tower and regularized in L3), the character $\eta(\tau)^{-1}$ we will meet in L4; (2.6) is a ratio of theta and eta functions in disguise. Keep both in view: the boson will become the Heisenberg vertex algebra and the fermion a Clifford (super) vertex algebra in L5.

2.3 The "partition function" pun is literal

Expand (2.5) as a power series in q:

$$\prod_{n\geq 1} \frac{1}{1-q^n} = \sum_{N\geq 0} p(N)\, q^N = 1 + q + 2q^2 + 3q^3 + 5q^4 + 7q^5 + \cdots, \tag{2.7}$$

where $p(N)$ is the number of ways of writing N as a sum of positive integers—the number of *integer partitions* of N. Euler's generating function and the physicist's partition function of a bosonic tower are *the same series*. (The coincidence of names is a genuine coincidence, but a happy one.) Likewise the fermionic product (2.6) counts partitions of N into *distinct* parts:

$$\prod_{n\geq 1} (1+q^n) = 1 + q + q^2 + 2q^3 + 2q^4 + 3q^5 + \cdots. \tag{2.8}$$

Computing a thermodynamic quantity and counting combinatorial objects are, here, one act.

2.4 Refining the trace: chemical potential and fugacity

Suppose the system also conserves a particle number N (an operator commuting with H). Holding $\langle N\rangle$ fixed *on average* introduces a second Lagrange multiplier, exactly as β was introduced for energy in Lecture 1. Call it the **chemical potential** μ. The weight becomes $\mathrm{e}^{-\beta(H-\mu N)}$, and the partition function is the grand canonical

$$\mathcal{Z} = \mathrm{Tr}\, \mathrm{e}^{-\beta(H-\mu N)} = \mathrm{Tr}\, z^N \mathrm{e}^{-\beta H}, \qquad z := \mathrm{e}^{\beta\mu}, \tag{2.9}$$

with z the **fugacity**. Physically μ is the energetic cost of adding one particle; formally z^N *grades* the trace by the conserved charge, and $\partial \log \mathcal{Z}/\partial(\beta\mu) = \langle N\rangle$.

One conjugate potential per conserved charge

This is the general move, stated once: for every commuting conserved charge Q_a we may insert a conjugate potential and refine the trace, $Z = \mathrm{Tr}\, \exp(-\sum_a \beta_a Q_a)$ (the *generalized Gibbs ensemble*). The sign-graded trace $\mathrm{Tr}\,(-1)^F \mathrm{e}^{-\beta H}$ is the same move with the fermion-number charge—the seed of the Witten index. Each such knob reappears, renamed, downstream: y^{J_0} on the torus (L4), the Cartan variables of a character (L5), and the y of the elliptic genus.

For the curious: the trace as a twisted character

The trace itself is the canonical map $\mathrm{End}(\mathcal{H}) \to \mathbb{C}$ for a dualizable object $\mathcal{H}$—the categorical fact the Finale will make precise. The graded trace $\mathrm{Tr}\, g\, \mathrm{e}^{-\beta H}$ for a symmetry operator g (here $g = z^N$ or $g = (-1)^F$) is a *twisted* version of that same map. When g generates a finite or compact group, summing or integrating over g projects onto invariants—the orbifold/gauging construction. We will not need orbifolds, but it is worth knowing that "insert an operator in the trace" is the universal hook on which charges, boundary conditions, and (in L4) spin structures all hang.

Rosetta notes

- Trace $\leftrightarrow$ sum over states in any basis, now basis-independent.
- $\mathrm{e}^{-\beta H}$ $\leftrightarrow$ evolution by imaginary time β — the cliffhanger for L3.
- Graded trace $\mathrm{Tr}\, z^N \mathrm{e}^{-\beta H}$ $\leftrightarrow$ inserting a symmetry operator $\leftrightarrow$ choosing a boundary condition $\leftrightarrow$ a spin structure (L3–L4).

Interlude: the Boltzmann factor as a character

For the reader from either camp

This chapter has no new computations. It names the single idea that the rest of the notes re-spell in five notations. If you read only one chapter slowly, read this one.

We now have both incarnations of the Boltzmann factor on the table: the *number* $\mathrm{e}^{-\beta E_i}$ of Lecture 1 and the *operator* $\mathrm{e}^{-\beta H}$ of Lecture 2. Before the object fragments into the notations of field theory, conformal field theory, and representation theory, let us abstract it. This is the piece that most often falls through the cracks: one learns each incarnation locally and never the idea behind all of them.

I.1 Three faces of one object

Face 1: the max-entropy weight. Maximizing entropy subject to fixed mean values of conserved charges Q_a forces the weight

$$w = \exp\Big(-\sum_a \beta_a Q_a\Big), \tag{I.1}$$

the β_a being the Lagrange multipliers (Lecture 1, generalized). Nothing else extremizes entropy at fixed charges. This is why the exponential is not a modelling choice but a theorem: the Boltzmann factor *is* the max-entropy weight.

Face 2: the character (a homomorphism). The map

$$Q \longmapsto \mathrm{e}^{-\beta Q} \tag{I.2}$$

is a continuous homomorphism from the *additive* monoid of conserved charges to the *multiplicative* monoid of weights—and, up to the choice of the constant β, the only one:

$$\mathrm{e}^{-\beta(Q+Q')} = \mathrm{e}^{-\beta Q}\,\mathrm{e}^{-\beta Q'}. \tag{I.3}$$

Additivity of charges becomes multiplicativity of weights. This single fact is why independent systems' partition functions multiply (§1.4), why Z factorizes over modes, and why product formulas—$\prod(1-q^n)^{-1}$, then η, then θ—pervade the subject. A one-dimensional character is the embryo of every "character" in Lectures 4 and 5.

Face 3: Euclidean evolution. Real-time evolution is the unitary $\mathrm{e}^{-\mathrm{i}Ht}$. Under the Wick rotation $t \to -\mathrm{i}\beta$ its oscillating phase becomes the real, decaying weight $\mathrm{e}^{-\beta H}$: through this analytic continuation inverse temperature is identified with an interval of imaginary (Euclidean) time, and β measures its *extent*—a duration, not a temperature. Taking the trace glues the two ends of that interval into a circle of circumference β (L3). In a two-dimensional CFT the space is itself a circle—the theory on a cylinder—so there are already *two* periodic directions, space and Euclidean time, and closing the time circle by the trace turns the cylinder into a torus (L4).

Remark. These are not three analogies; they are three theorems about the same function. Face 1 says *which* function; Face 2 says *how it composes*; Face 3 says *what geometry it secretly carries*.

I.2 The general Boltzmann factor

Collect every commuting conserved charge—energy H, particle number N, momentum, an internal or flavour charge J_a—and pair each with a **conjugate potential**:

$$\boxed{w = \exp\Big(- \beta H + \beta\mu N + \textstyle\sum_a \nu_a J_a + \cdots \Big), \qquad Z = \mathrm{Tr}\ w.} \tag{I.4}$$

Two structural facts organize (I.4):

- **Extensive charge $\leftrightarrow$ intensive potential** is a Legendre duality: the pairs (E, β), (N, μ), (J_a, ν_a). Differentiating $\log Z$ by a potential returns the mean of its charge; the second derivative returns the fluctuations.
- The **chemical potential** is the prototype "extra knob." It weights states by how much conserved charge they carry; equivalently a fugacity $z = \mathrm{e}^{\beta\mu}$ grades the trace. *Refining the trace is turning a knob.*

To see a knob actually move the grading, watch **spin**. A conserved angular-momentum projection J has quantized eigenvalues: whole units $\ldots, -1, 0, 1, \ldots$ on integer-spin (bosonic) states, half-units $\ldots, -\frac{3}{2}, -\frac{1}{2}, \frac{1}{2}, \ldots$ on half-integer-spin (fermionic) ones. Its conjugate potential is a rotation angle θ, which enters the trace as a *holonomy*—a twist applied once around the Euclidean time circle of Face 3—so that a state of spin j picks up the phase $\mathrm{e}^{\mathrm{i}\theta j}$ on the way round:

$$\mathrm{Tr}\ \mathrm{e}^{\mathrm{i}\theta J}\,\mathrm{e}^{-\beta H} \ = \ \sum_j \mathrm{e}^{\mathrm{i}\theta j}\ \mathrm{Tr}_{\mathcal{H}_j}\,\mathrm{e}^{-\beta H},$$

one term for each spin sector $\mathcal{H}_j$. Although θ runs continuously, nothing depends on it except through that holonomy—and what the holonomy detects is the *quantization* of j. At the trivial twist $\theta = 0$ every sector is counted equally; at a full turn $\theta = 2\pi$ the phase becomes $\mathrm{e}^{2\pi \mathrm{i} J} = (-1)^F$, the familiar sign a fermion picks up under a 2π rotation—the coarse $(-1)^F$ grading of L3–L4 in embryo.

Particle number and energy grade the trace the same way; only the *spectrum* being sorted changes. Particle number is likewise quantized, so its fugacity $z = \mathrm{e}^{\beta\mu}$ organizes the trace into the grand-canonical series $\sum_N z^N Z_N$, and raising μ makes each added particle cheaper, tilting the sum toward larger N. Energy is not quantized in integer units: β weights each state by its actual energy through the smooth envelope $\mathrm{e}^{-\beta E}$, sliding the weight rather than sorting it—raising β (cooling) suppresses high-energy states toward the ground state, lowering it spreads the weight across the whole range. In every case the exponent is a charge and the knob its conjugate potential; whether turning it sorts the trace into discrete sectors or slides a continuous weight is fixed by the charge, not the knob.

An aside on scale. Because β is the circumference of the Euclidean time circle (Face 3), it is also a length; dialing it from small to large carries the theory from short-distance toward long-distance physics. Turning the β knob is, in this sense, moving along a renormalization-group scale—a connection we flag here and leave to a course on critical phenomena.

I.3 The same knobs, renamed, downstream

Here is the cross-disciplinary payoff to keep visible for the rest of the notes. Each conjugate potential we meet now reappears later wearing a geometric or formal name.

Conjugate potential	**in L1–L2 it is**	**downstream it becomes**
β (to energy H)	inverse temperature	$\operatorname{Im}\tau$ on the torus (L4); the modular parameter of a character (L5)
μ (to charge N)	chemical potential / fugacity z	the *elliptic variable* y conjugate to J_0 (L4); Cartan variables of Weyl–Kac (L5); the y of the elliptic genus
θ (to fermion number F)	sign grading: $(-1)^F$ at $\theta = 2\pi$	a spin structure (L3–L4); the Witten index / elliptic genus (L5)
source (to a current)	external field	a background gauge field / Wilson line (L3–L4)

For the curious: a categorical wink

Face 2 says a Boltzmann factor *is* a monoidal functor $(\text{charges}, +, 0) \to (\mathbb{C}, \times, 1)$. That is the one-object baby version of the bordism functor $Z\colon \mathrm{Bord}_2 \to \mathrm{Vect}$ of the finale: assign weights compatibly with a composition law. From this height, "partition function = trace = character" is the single statement that the assignment is *functorial*.

A checklist for everything that follows

Whenever a later chapter writes a new $\mathrm{e}^{-(\cdots)}$ or $q^{(\cdots)}$, ask three questions:

1. *Which charges sit in the exponent?*
2. *What are their conjugate potentials?*
3. *What geometric object closes the trace?*

Those three answers are the entire dictionary we are about to assemble.

Chapter 3

Quantum Field Theory: the partition function as a path integral

For the reader from mathematics

The path integral is, for our purposes, a (formal, then regularized) Gaussian measure. The one result you need to trust is that a Gaussian integral produces a determinant, and that for fields on a circle that determinant is computed by a regularized product over modes. The new geometric input is that the trace of Lecture 2 *is* a path integral on a circle, whose circumference is β.

3.1 The trace is a path integral on a circle

Face 3 of the Interlude recast $\mathrm{e}^{-\beta H}$ as evolution through an interval β of imaginary time. To make a path integral of it we need a basis of states to propagate between, and the natural one is the basis of **field-configuration eigenstates**. Just as a single particle has position eigenstates $\hat{x}|x\rangle = x|x\rangle$, a field has eigenstates of the (Schrödinger-picture) field operator,

$$\hat{\phi}(\mathbf{x})|\phi\rangle = \phi(\mathbf{x})|\phi\rangle, \qquad \hat{\phi}(\mathbf{x}) = \sum_k \frac{1}{\sqrt{2\omega_k}}\big(a_k\, u_k(\mathbf{x}) + a_k^\dagger\, u_k^*(\mathbf{x})\big),$$

one for every classical profile $\phi(\mathbf{x})$, built from the mode functions u_k and the ladder operators $a_k, a_k^\dagger$ of Lecture 2; they resolve the identity, $\int \mathcal{D}\phi\, |\phi\rangle\langle\phi| = \mathbf{1}$.

These $|\phi\rangle$ are vectors in the very Fock space assembled in §2.2—the symmetric or exterior algebra on the modes. That space is *graded*: it splits into subspaces of definite energy, $\mathcal{H} = \bigoplus_n \mathcal{H}_n$ (and, independently, into the particle-number sectors carried by the fugacity z^N of §2.4). A number eigenstate lives in a single grade; a field eigenstate $|\phi\rangle$ is a superposition spread across all of them, as a position eigenstate spreads over the levels of an oscillator. Keep the word *grading* in view: this same energy grading returns as the q^{L_0} of the torus (L4) and the graded dimension of a character (L5).

Now compute the trace by sandwiching, exactly as in §2.2 but in the field basis. Split $\mathrm{e}^{-\beta H} = (\mathrm{e}^{-\epsilon H})^N$ into N steps of size $\epsilon = \beta/N$ and insert the resolution of the identity between every pair:

$$Z = \mathrm{Tr}\ \mathrm{e}^{-\beta H} = \int \mathcal{D}\phi\, \langle\phi|\mathrm{e}^{-\beta H}|\phi\rangle = \int \prod_{i=1}^{N} \mathcal{D}\phi_i\ \prod_{i=1}^{N} \langle\phi_{i+1}|\mathrm{e}^{-\epsilon H}|\phi_i\rangle, \qquad \phi_{N+1} \equiv \phi_1.$$

The trace has already done its work in setting $\phi_{N+1} = \phi_1$: the chain of configurations closes into a loop. Each infinitesimal factor is the Boltzmann weight of a single imaginary-time step,

$$\langle\phi_{i+1}|\mathrm{e}^{-\epsilon H}|\phi_i\rangle \ \simeq\ \exp\big(-\epsilon\, L_E[\phi_i, \dot{\phi}_i]\big), \qquad \dot{\phi}_i = \tfrac{1}{\epsilon}(\phi_{i+1} - \phi_i),$$

with L_E the Euclidean Lagrangian. Here is the rhyme with the Interlude: the exponent is *additive* over the steps, so the weights *multiply*—the very homomorphism of Face 2 (additivity $\to$ multiplicativity) that made Z factorize over independent systems, now applied to successive slices of imaginary time. Their product assembles the exponential of the action, $\prod_i \mathrm{e}^{-\epsilon L_E} = \mathrm{e}^{-\sum_i \epsilon L_E} \to \mathrm{e}^{-\int_0^\beta \mathrm{d}\tau\, L_E} = \mathrm{e}^{-S_E[\phi]}$, and the sum over the intermediate configurations ϕ_i becomes an integral over field histories $\phi(\tau, \mathbf{x})$, their two ends glued by the trace into a circle:

$$\boxed{Z = \mathrm{Tr}\ \mathrm{e}^{-\beta H} = \int_{\phi(\tau+\beta)=\pm\phi(\tau)} \mathcal{D}\phi\ \mathrm{e}^{-S_E[\phi]}.} \tag{3.1}$$

The thermal partition function is a path integral over fields on a Euclidean time circle S^1_β of circumference β. *Inverse temperature has become a length.*

The same construction runs for a fermion, with one twist. Because its ladder operators *anticommute* ($\{c, c^\dagger\} = 1$, Lecture 2), the useful basis is the fermionic **coherent states** $|\psi\rangle$, eigenstates of the annihilator $c|\psi\rangle = \psi|\psi\rangle$ whose labels ψ are anticommuting **Grassmann** numbers. Tracing over them carries an extra sign,

$$\mathrm{Tr}\ A = \int \mathrm{d}\bar{\psi}\, \mathrm{d}\psi\ \mathrm{e}^{-\bar{\psi}\psi}\, \langle -\psi|A|\psi\rangle,$$

and it is the $\langle -\psi|$—the fermionic trace closing on *minus* the starting configuration—that will force the antiperiodic boundary condition of the next section.

3.2 Statistics dictates boundary conditions

What is the $\pm$ in (3.1)? Tracing means gluing $\tau = 0$ to $\tau = \beta$, and the gluing condition is fixed by the statistics of the field:

$$\text{bosons:}\quad \phi(\tau+\beta) = +\phi(\tau) \qquad \text{(periodic)}, \tag{3.2}$$

$$\text{fermions:}\quad \psi(\tau+\beta) = -\psi(\tau) \qquad \text{(antiperiodic)}. \tag{3.3}$$

The minus sign for fermions is exactly the $(-1)^F$ of §2.4: the trace of $\mathrm{e}^{-\beta H}$ over a fermionic mode equals the antiperiodic path integral, while inserting $(-1)^F$ would flip it to periodic. This is the path-integral form of the Kubo–Martin–Schwinger (KMS) condition, whose proper home—correlation functions—is the business of the box below.

Before refining further, pause on what kind of choice this $\pm$ is. A fermion is a spinor, and to put spinors on a spacetime at all one must equip it with a **spin structure**: a consistent rule for the sign a spinor picks up when carried around each closed loop. In vacuum QFT on $\mathbb{R}^d$ this never comes up, and for good reason—$\mathbb{R}^d$ is simply connected, its spin structure is unique, and there is no choice to make. But taking the trace changed the topology under our feet: thermal spacetime is $\mathsf{S}^1_\beta \times \mathbb{R}^{d-1}$, and a circle admits exactly *two* spin structures—the two gluings $\psi(\tau + \beta) = \pm\psi(\tau)$ just written. The partition function cannot be agnostic between them: Tr $\mathrm{e}^{-\beta H}$ *is* the antiperiodic structure, and inserting $(-1)^F$ switches to the periodic one. "Statistics dictates boundary conditions" is, said geometrically, "taking a trace selects a spin structure."

Nor is this exotic; it has been hiding in plain sight in every finite-temperature calculation. Fourier analysis on the time circle assigns a periodic (bosonic) field the **Matsubara frequencies** $\omega_n = 2\pi n/\beta$, but an antiperiodic (fermionic) field the *odd* ones, $\omega_n = (2n + 1)\pi/\beta$. Those odd integers, familiar from any first pass through thermal field theory, are a choice of spin structure in action.

For the curious: equilibrium is what makes imaginary time a circle

Why should *equilibrium*, of all things, hand every field a discrete comb of frequencies? Lecture 1 characterized the thermal state variationally: $\mathrm{e}^{-\beta H}/Z$ is the unique entropy maximizer at fixed mean energy (Face 1 of the Interlude). The KMS condition characterizes the same state geometrically: if the correlation functions of a finite system continue to imaginary time and close up (anti)periodically with period β, the state is forced to be $\mathrm{e}^{-\beta H}/Z$—cyclicity of the trace, read backwards (Face 3). The two characterizations select the same state, and the geometric one is the one that survives the thermodynamic limit: in the algebraic formulation of Haag, Hugenholtz and Winnink, a KMS state is what "equilibrium at inverse temperature β" *means* for an infinite system, where the Gibbs sum may no longer exist. The Matsubara comb is then the spectral fingerprint of equilibrium: a correlator periodic on a circle of circumference β carries Fourier support only at $\omega_n = 2\pi n/\beta$ (odd multiples of π/β in the antiperiodic case), while out of equilibrium nothing glues imaginary time into a circle and the frequency content spreads off the comb. The Lagrange multiplier of Lecture 1 has become a geometric period: maximizing entropy and compactifying Euclidean time are the same act, and the comb spacing $2\pi/\beta$ is the entropy multiplier turned into a frequency quantum.

That $\pm$ is only the coarsest twist. In §I.2 we refined the trace by a conjugate potential $\mathrm{e}^{\mathrm{i}\theta Q}$ for any conserved charge Q; in the path integral each such insertion becomes a *twisted boundary condition*—the field returns to itself only up to the symmetry group element $g = \mathrm{e}^{\mathrm{i}\theta Q}$ that the insertion generates,

$$\mathrm{Tr}\left(\mathrm{e}^{\mathrm{i}\theta Q}\,\mathrm{e}^{-\beta H}\right) = \int_{\phi(\tau+\beta)=\mathrm{e}^{\mathrm{i}\theta}\phi(\tau)} \mathcal{D}\phi\;\mathrm{e}^{-S_E[\phi]} \qquad \text{(unit charge).} \tag{3.4}$$

The twist is precisely a flat *background* connection for the global symmetry generated by Q: a holonomy g around S^1_β, switched on by hand and not summed over. What it looks like depends on the group Q generates.

If that group is *finite*—fermion parity $(-1)^F$ generating $\mathbb{Z}_2$ is the case in hand—the holonomy g takes only finitely many values, each labelling a distinct *twisted sector*, with no continuum in between: antiperiodicity is the nontrivial element of $\mathbb{Z}_2$ and periodicity the identity. *Summing* over the sectors, instead of fixing one, gauges the symmetry—the orbifold construction.[1]

If instead Q generates a *continuous* group such as $U(1)$, the holonomy $\mathrm{e}^{\mathrm{i}\theta}$ runs continuously and the twist is a background $U(1)$ gauge field wrapped on the circle. This is the chemical potential exactly: μ enters as an imaginary time-component gauge field through $\partial_\tau \to \partial_\tau - \mu$, a **Wilson line** $\mathrm{e}^{\beta\mu}$ around S^1_β, and one *dials* it rather than summing. Either way Q must generate a *global* symmetry, acting on the states the trace runs over, and the connection is a *background*: a gauged (local) symmetry grades nothing—its charged states are projected out—and its connections are integrated over, not fixed. The holonomy θ of §I.2 is now literal, the phase a field acquires on going once around S^1_β—every knob of the Interlude a boundary condition.

Two choices now, four spin structures later

A field on the time circle carries one of two boundary conditions (periodic "R" or antiperiodic "NS"). On the torus of Lecture 4 there are *two* circles, hence $2 \times 2 = 4$ boundary-condition choices: the four spin structures. The single sign we are choosing here is one bit of that data.

[1] More precisely, a twist is a flat G-bundle on the spacetime: on S^1_β a holonomy $g \in G$ (a conjugacy class), on the torus of Lecture 4 a commuting pair (g, h) of holonomies for the two cycles. When one sums over sectors to gauge G, the phases that may be attached to them form a torsor over $H^2(G, U(1))$—Vafa's *discrete torsion*—while the obstruction to gauging at all is a 't Hooft anomaly in $H^3(G, U(1))$.

3.3 Free fields, exactly

For a free (Gaussian) theory the path integral is computed exactly by a functional determinant,

$$Z = \big(\det \mathcal{O}\big)^{\mp 1/2}, \qquad \begin{cases} -\frac{1}{2} & \text{real (scalar) boson } (\mathcal{O} = -\partial^2 + m^2), \\ +\frac{1}{2} & \text{Majorana fermion } (\mathcal{O} = \not{\partial} + m). \end{cases} \tag{3.5}$$

The opposite signs of the exponent are the determinant-level shadow of Bose vs. Fermi, and the halves are the shadow of *reality*: a real (Majorana) fermion's Gaussian integral is a Pfaffian, $\mathrm{Pf}\,\mathcal{O} = \det^{1/2}\mathcal{O}$, just as the real scalar gives $\det^{-1/2}$; their complex cousins—Dirac fermion, complex scalar—carry the powers ± 1. Writing the determinant as a product over the modes on S^1_β and doing the Matsubara sum reproduces (2.5) and (2.6), along with the zero-point factor the next section takes up: the path integral and the trace agree, as they must. Meanwhile $\log Z$ is the sum of connected vacuum diagrams—the field-theory free energy.

3.4 The vacuum energy and the first $-1/12$

The determinant carries a zero-point piece $\frac{1}{2}\sum_n \omega_n$, the sum of the ground-state energies of all the modes (Example 1.2, summed). It diverges. Make it finite the cheap way: put the system in a periodic box of length L, an infrared regulator (a lattice length, nothing geometric yet), so a tower of modes discretizes to $\omega_n = 2\pi n/L$ and the zero-point sum is $\frac{\pi}{L}\sum_{n\geq 1} n$. (Reading L as a genuine spatial circle—one cycle of a torus—is the business of the 2d conformal field theory of Lecture 4; here it is only a cutoff.) Zeta-regularize:

$$\sum_{n\geq 1} n \overset{\zeta}{=} \zeta(-1) = -\frac{1}{12}, \qquad \Longrightarrow \qquad E_0 = -\frac{\pi}{12L}. \tag{3.6}$$

The finite leftover is the **Casimir energy**.[2]

What does the fermion contribute? Its sign is fixed by the mode algebra, before any geometry enters. Quantizing a bosonic mode symmetrizes the classical product: $H = \frac{\omega}{2}(a^\dagger a + a\,a^\dagger) = \omega(a^\dagger a + \frac{1}{2})$, the $+\frac{1}{2}\omega$ paid to the commutator $[a, a^\dagger] = 1$. Quantizing a fermionic mode antisymmetrizes: $H = \frac{\omega}{2}(c^\dagger c - c\,c^\dagger) = \omega(c^\dagger c - \frac{1}{2})$, and now the anticommutator $\{c, c^\dagger\} = 1$ pays $-\frac{1}{2}\omega$. The same one bit of statistics that set the $\pm$ of §3.2 flips the zero-point energy: every bosonic mode contributes $+\frac{1}{2}\omega$ to the vacuum, every fermionic mode $-\frac{1}{2}\omega$.

That opposite sign is an invitation. Suppose a theory supplies one fermionic mode for every bosonic mode of the same frequency—matched mode for mode, not component for component. (The distinction matters: a fermion's action is first order, so the field is its own conjugate momentum and each component carries only *half* a degree of freedom; a naive component count makes fermions look twice as numerous as they are. Which multiplets match is then dimension-dependent bookkeeping—in four dimensions, for instance, the four real components of a Majorana fermion supply the same two modes per momentum as a complex scalar.) Then $+\frac{1}{2}\omega$ meets $-\frac{1}{2}\omega$ at every single frequency and the vacuum energy cancels mode by mode: divergence, finite part, everything, with no regulator invoked at all. This is historically one of the cheapest ways to motivate **supersymmetry**: a symmetry exchanging bosons and fermions forces exactly such a pairing of modes, and a vanishing vacuum energy is its signature—the SUSY algebra makes $H \geq 0$, saturated by an invariant ground state. What survives when the pairing is imperfect—the net count of unpaired ground states—is the **Witten index** of Lecture 5, a partition function that counts with signs.

When the spectra do *not* match, the regulator exposes the mismatch—and here the box itself asks the fermion the same question the time circle asked in §3.2: periodic or antiperiodic

[2] The textbook Casimir energy of a 2d field on a circle, $-\pi c/6L$, counts the left- and right-moving towers together; our single tower $\omega_n = 2\pi n/L$ is one chiral half of it.

around the compact direction? Both are consistent; as with the two spin structures there, the regulated theory must simply pick one. (On the torus of Lecture 4 the two choices become the R and NS sectors.) Periodic, the fermion is integer moded like the boson, and its zero-point sum is the boson's with the sign reversed:

$$E_0^{\mathrm{per}} = -\tfrac{1}{2}\sum_{n\geq 1}\omega_n = -\frac{\pi}{L}\sum_{n\geq 1} n \overset{\zeta}{=} +\frac{\pi}{12L}, \tag{3.7}$$

equal and opposite to (3.6)—the finite shadow of the mode-by-mode cancellation above. Antiperiodic, it is half-integer moded, $\omega_r = 2\pi r/L$ with $r = \frac{1}{2}, \frac{3}{2}, \dots$, and

$$E_0^{\mathrm{anti}} = -\tfrac{1}{2}\sum_{r}\omega_r = -\frac{\pi}{L}\sum_{k\geq 0}\left(k+\tfrac{1}{2}\right) \overset{\zeta}{=} -\frac{\pi}{24L}, \tag{3.8}$$

the shifted sum $\sum_{k\geq 0}(k+\frac{1}{2})$ ζ-regularizing to $+\frac{1}{24}$. Two boundary conditions, two vacuum energies. Notice, finally, what did *not* spoil the cancellation: Euclidean time. The zero-point sum is a statement about the Hamiltonian on space, untouched by the Wick rotation, and the thermal antiperiodicity of §3.2 lives on the *time* circle, which never enters it. What halved the antiperiodic answer—$-\pi/24L$ against the boson's $-\pi/12L$—is the shifted moding alone, $\sum(k+\frac{1}{2})$ against $\sum n$. On the torus these two numbers reappear as the leading powers of the theta-function blocks of Lecture 4.

The thread, continued

The $\zeta(-1) = -\frac{1}{12}$ of (3.6) is the regularized shadow of the zero-point factor $\mathrm{e}^{-\beta\hbar\omega/2}$ of Lecture 1, summed over the whole tower. Once L is a genuine circle (Lecture 4), the coefficient is the central charge: the single tower carries $E_0 = -\frac{c}{24}\cdot\frac{2\pi}{L}$ per chiral half, with $c = 1$ for the free boson—the $q^{1/24}$ inside $\eta(\tau)$. Left- and right-movers together give $2\times(-\frac{1}{24}) = -\frac{1}{12}$. In Lecture 5 the same $-c/24$ is the universal grading shift $L_0 - c/24$.

For the curious: why ζ-regularization is legitimate

$\zeta(-1) = -\frac{1}{12}$ is not a claim that $1+2+3+\cdots$ converges. It is the statement that the analytic continuation of $\zeta(s) = \sum n^{-s}$ (defined for $\Re s > 1$) takes the value $-\frac{1}{12}$ at $s=-1$. That value is *not* a convention: by the identity theorem a holomorphic function admits *at most one* analytic continuation to a given connected domain, so ζ's extension—and with it the number $-\frac{1}{12}$—is forced the moment one insists on analyticity. This uniqueness is why the continuation is the scheme-independent finite part of the regularized vacuum energy: different regulators (heat kernel, point-splitting) must agree on it, the divergent pieces being local and removed by counterterms. The physical Casimir force, measured in the laboratory in its three-dimensional electromagnetic version, confirms that such finite remainders are real physics.

Rosetta notes

- $\beta \leftrightarrow$ length of the Euclidean time circle.
- $(-1)^F$ insertion $\leftrightarrow$ flipping the fermion from antiperiodic to periodic $\leftrightarrow$ the choice of spin structure on S^1_β (§3.2), one bit of the four of L4.
- det $\leftrightarrow$ the one-loop free energy; its zero-point part $\leftrightarrow$ the $-c/24$ shift.

For the curious: gluing is composition (a categorical preview)

Sewing two path integrals along a shared boundary slice composes the corresponding linear maps ("propagators"). The trace of Lecture 2 is the special sewing that closes a cylinder into a circle. A partition function on a closed manifold is then the value of a functor on a cobordism $\varnothing \to \varnothing$—a number. We make this precise in the finale; for now it explains why (3.1) "closing the interval into a circle" is the geometric form of taking a trace.

Chapter 4

Two-dimensional CFT: the partition function on the torus

For the reader from mathematics

The "physics torus" is your elliptic curve $\mathbb{C}/(\mathbb{Z}+\tau\mathbb{Z})$, and "modular invariance" is the action of $SL(2,\mathbb{Z})$ on $\tau \in \mathbb{H}$. The partition function will turn out to be built from the Dedekind η and the Jacobi θ functions—weight-$\frac{1}{2}$ modular forms you already know. What CFT adds is a *reason* those particular forms appear: they are traces over representations.

For the reader from physics

The one structural novelty is that the time circle of Lecture 3 is now accompanied by a *space* circle, and we are free to relabel which is which. The partition function cannot depend on that relabelling—that constraint is modular invariance, and it is as strong as it is because $SL(2,\mathbb{Z})$ is large.

4.1 Two circles make a torus

Put a 2d CFT on a spatial circle of circumference 2π. Its states live in a Hilbert space $\mathcal{H}$ acted on by two copies of the Virasoro algebra—generators L_n for the left-movers, $\bar{L}_n$ for the right—and the zero modes do most of this chapter's work. The one statement to anchor everything, especially if the lingo is new, is this: *the Hamiltonian of the theory on the circle is*

$$H \;=\; L_0 + \bar{L}_0 - \frac{c}{12}, \qquad\qquad P \;=\; L_0 - \bar{L}_0,$$

with P the momentum, generating translation around the circle. L_0 and $\bar{L}_0$ are nothing exotic: they are the left- and right-moving halves of the energy. And the shift $-c/12$ is the Casimir energy of Lecture 3, one $-c/24$ for each half. Evolving in Euclidean time by β and tracing makes a torus: one cycle is the space circle, the other the time circle. Allowing the time evolution to also *translate* along the space circle by an angle θ (a "twist") combines the two real parameters into one complex **modular parameter** $\tau = \tau_1 + \mathrm{i}\tau_2 \in \mathbb{H}$, with $\beta = 2\pi\tau_2$ and $\theta = 2\pi\tau_1$, and we set $q = \mathrm{e}^{2\pi\mathrm{i}\tau}$. The literature calls this q the **nome** of the torus—an old word from elliptic-function theory, worth recognizing but nothing mysterious: it is the natural expansion variable for anything living on the torus, and, physically, a Boltzmann weight once more. Indeed $|q| = \mathrm{e}^{-\beta}$, so q^{L_0} suppresses a state by $\mathrm{e}^{-\beta}$ and rotates its phase by θ per unit of L_0. **Mind that it is not the thermal $q = \mathrm{e}^{-\beta\hbar\omega}$ of L2**, which weighted the quanta of one chosen oscillator; the two coincide only once the tower is graded by L_0 in integer steps. The torus is the quotient $\mathbb{C}/(\mathbb{Z}+\tau\mathbb{Z})$.

Why *two* copies of the Virasoro algebra, and why the constant talk of holomorphy? In two Euclidean dimensions, conformal geometry and complex analysis are the same subject: pack the two worldsheet coordinates into one complex z, and a map preserves angles precisely where it is holomorphic (or antiholomorphic)—the Cauchy–Riemann equations *are* the conformality condition. The 2d conformal "group" is therefore infinite-dimensional: every holomorphic $z \mapsto z + \epsilon(z)$ is an infinitesimal symmetry, one generator $\ell_n = -z^{n+1}\partial_z$ per Fourier mode of ϵ. Quantum mechanically their algebra acquires a central term, measured by the same c that set the Casimir energy a paragraph ago:

$$[L_m, L_n] = (m-n)\, L_{m+n} + \frac{c}{12}\, m(m^2-1)\, \delta_{m+n,0}, \tag{4.1}$$

one copy for the holomorphic side and an independent copy $\bar{L}_n$ for the antiholomorphic. Everything in the theory accordingly splits into left-movers (functions of z) and right-movers (functions of $\bar{z}$), coupled only through the spectrum—and the split survives the trace: a quantity built from left-movers alone depends on the modulus only through $q = \mathrm{e}^{2\pi \mathrm{i}\tau}$, never $\bar{q}$, so *chiral* objects will be holomorphic functions of τ. That innocuous remark is the engine of §4.6.

The partition function is the graded trace

$$\boxed{Z(\tau, \bar{\tau}) = \mathrm{Tr}_{\mathcal{H}}\; q^{L_0 - \frac{c}{24}}\; \bar{q}^{\bar{L}_0 - \frac{c}{24}}.} \tag{4.2}$$

Compare (2.1)—and note this is an identity, not an analogy: unpacking $q = \mathrm{e}^{2\pi \mathrm{i}\tau}$,

$$q^{L_0 - \frac{c}{24}}\, \bar{q}^{\bar{L}_0 - \frac{c}{24}} \;=\; \mathrm{e}^{-\beta H}\, \mathrm{e}^{\mathrm{i}\theta P},$$

so (4.2) *is* the thermal trace of Lecture 2, refined by a conjugate potential for the momentum. Inverse temperature has been complexified into τ.

Where the $-c/24$ lands

The shift $L_0 \to L_0 - \frac{c}{24}$ in (4.2) is exactly the Casimir energy of Lecture 3: putting the theory on a circle costs a ground-state energy $-c/24$ per chiral half. This is the same $-1/12$, now wearing its conformal clothes.

Reading the torus trace with the Interlude's checklist

Which charges? $H = L_0 + \bar{L}_0 - \frac{c}{12}$ and $P = L_0 - \bar{L}_0$. *Which conjugate potentials?* $\beta = 2\pi\tau_2$ and $\theta = 2\pi\tau_1$, packaged as the single complex τ. *What closes the trace?* The two cycles of the torus. If the theory has an extra current charge J_0, we may refine the trace to $\mathrm{Tr}\; y^{J_0} q^{L_0 - c/24} \cdots$: the new potential y is the *elliptic variable*—the chemical potential of §2.4, now a Jacobi-form variable.

4.2 Modular invariance

The same torus can be described by different choices of which lattice vectors are the two cycles. The relabellings form the modular group $SL(2,\mathbb{Z})$, acting by

$$\tau \longmapsto \frac{a\tau + b}{c\tau + d}, \qquad \begin{pmatrix} a & b \\ c & d \end{pmatrix} \in SL(2,\mathbb{Z}), \tag{4.3}$$

generated by

$$T:\; \tau \mapsto \tau + 1, \qquad S:\; \tau \mapsto -\frac{1}{\tau}. \tag{4.4}$$

Since they describe the *same* torus, the partition function of a consistent theory must be invariant:

$$Z(\tau+1) = Z(\tau), \qquad Z(-1/\tau) = Z(\tau). \tag{4.5}$$

T-invariance is a statement about the spectrum modulo the $-c/24$ shift; S-invariance exchanges the space and time circles, relating high and low "temperature." Modular invariance is the conformal descendant of the cyclicity of the trace (§2.1).

4.3 The free boson: the Dedekind eta

For a single non-compact free boson ($c = 1$) the trace (4.2) factorizes into the zero-mode (momentum) integral and the oscillator trace (2.5). The result, quoted per unit target-space volume (the momentum integral also supplies an overall volume factor), is

$$Z_{\text{boson}}(\tau) = \frac{1}{\sqrt{\operatorname{Im}\tau}\,|\eta(\tau)|^2}, \qquad \eta(\tau) = q^{1/24} \prod_{n\geq 1}(1-q^n). \tag{4.6}$$

There is the $-c/24$ thread again, as the $q^{1/24}$. The Dedekind eta is a modular form of weight $\frac{1}{2}$: $\eta(\tau+1) = e^{i\pi/12}\eta(\tau)$ and $\eta(-1/\tau) = \sqrt{-i\tau}\,\eta(\tau)$. The factor $1/\sqrt{\operatorname{Im}\tau}$ from the momentum integral is exactly what is needed to make (4.6) modular invariant. The oscillator product $\prod(1-q^n)^{-1}$ of Lecture 2 has become a modular object.

For the curious: the compact boson and T-duality

Compactify the boson on a circle of radius R. States now carry momentum n and winding w, and the zero-mode integral becomes a sum over the Narain lattice. The partition function is invariant under $R \to \alpha'/R$ with $(n,w) \to (w,n)$: **T-duality**. Tracking the Interlude's potentials, the S-transform of the torus and the $R \to 1/R$ duality are two faces of the same self-consistency—high/low temperature exchange.

4.4 The free fermion: theta functions and four spin structures

A free fermion on the torus must be assigned a boundary condition on *each* of the two cycles—periodic (Ramond, R) or antiperiodic (Neveu–Schwarz, NS)—by the rule of §3.2. That is $2 \times 2 = 4$ choices: the **four spin structures**. We write (a, b) with a the *space*-cycle and b the *time*-cycle condition; by §3.2, an R (periodic) time cycle means a $(-1)^F$ insertion in the trace. Each choice yields a trace that evaluates to a ratio of a Jacobi theta function and η:

$$(\text{NS, NS}) \sim \sqrt{\tfrac{\theta_3}{\eta}}, \quad (\text{NS, R}) \sim \sqrt{\tfrac{\theta_4}{\eta}}, \quad (\text{R, NS}) \sim \sqrt{\tfrac{\theta_2}{\eta}}, \quad (\text{R, R}) \sim \sqrt{\tfrac{\theta_1}{\eta}} \equiv 0, \tag{4.7}$$

the last vanishing because $\theta_1 \equiv 0$ (a fermion zero mode; the same zero mode leaves the R ground states degenerate). Parallel to η in (4.6), let us spell the **Jacobi theta constants** out as products,

$$\theta_3 = \prod_{n\geq 1}(1-q^n)(1+q^{n-\frac{1}{2}})^2, \quad \theta_4 = \prod_{n\geq 1}(1-q^n)(1-q^{n-\frac{1}{2}})^2, \quad \theta_2 = 2\,q^{1/8}\prod_{n\geq 1}(1-q^n)(1+q^n)^2 \tag{4.8}$$

(equivalently the sums $\sum_{n\in\mathbb{Z}} q^{n^2/2}$, $\sum_{n\in\mathbb{Z}}(-1)^n q^{n^2/2}$ and $\sum_{n\in\mathbb{Z}} q^{(n+1/2)^2/2}$). The blocks are then the Lecture 2 products in thin disguise:

$$\sqrt{\tfrac{\theta_3}{\eta}} = q^{-\frac{1}{48}} \prod_{r=\frac{1}{2},\frac{3}{2},\ldots}(1+q^r), \qquad \sqrt{\tfrac{\theta_4}{\eta}} = q^{-\frac{1}{48}} \prod_{r=\frac{1}{2},\frac{3}{2},\ldots}(1-q^r), \qquad \sqrt{\tfrac{\theta_2}{\eta}} = \sqrt{2}\,q^{\frac{1}{24}} \prod_{n\geq 1}(1+q^n)\,: \tag{4.9}$$

the fermionic partition function (2.6), mode by mode, with half-integer (NS) or integer (R) moding, a sign when $(-1)^F$ is inserted, a $\sqrt{2}$ for the R zero mode, and the sector ground-state energy out front. The leading powers in (4.9) are old friends: $q^{-1/48}$ for NS, $q^{+1/24}$ for R—the two regulator-box vacuum energies (3.8) and (3.7) of Lecture 3, in units of $2\pi/L$, now reading as $-\frac{c}{24}$ and $h-\frac{c}{24}$ for the $c=\frac{1}{2}$ Majorana fermion. The NS ground state is the true vacuum; the R ground state is *not*—it is a state of dimension $h=\frac{1}{16}$, created by the *spin field* of the Ising model, sitting at $\frac{1}{16}-\frac{1}{48}=+\frac{1}{24}$. The boundary conditions of Lecture 3 have become sector data, each sector with its own ground-state energy. Under the modular group, T and S *permute* the spin structures, hence permute $\theta_2, \theta_3, \theta_4$:

$$T:\ \theta_3 \leftrightarrow \theta_4,\quad \theta_2 \to \theta_2; \qquad S:\ \theta_2 \leftrightarrow \theta_4,\quad \theta_3 \to \theta_3. \tag{4.10}$$

A modular-invariant fermion partition function is therefore a specific sum over spin structures, $\frac{1}{2}(|\theta_3/\eta| + |\theta_4/\eta| + |\theta_2/\eta|)$ in schematic form—the content of the 2d Ising model at its critical point, the $c=\frac{1}{2}$ theory realized by a single Majorana fermion.

One more enlargement before the machinery of Lecture 5. Put the boson and the fermion *together* ($c = 1+\frac{1}{2}=\frac{3}{2}$), and the combined theory carries more symmetry than two commuting Virasoros: alongside the stress tensor there is a fermionic current $G(z)$ of weight $\frac{3}{2}$—concretely $G \sim \psi\,\partial\phi$—whose modes G_r extend (4.1) to the $\mathcal{N}=1$ **super-Virasoro algebra**,

$$[L_m, G_r] = \left(\tfrac{m}{2}-r\right) G_{m+r}, \qquad \{G_r, G_s\} = 2L_{r+s} + \tfrac{c}{3}\left(r^2-\tfrac{1}{4}\right)\delta_{r+s,0}. \tag{4.11}$$

This is **superconformal** symmetry: conformal transformations together with supersymmetries rotating the boson into the fermion, holomorphic like everything else in this chapter.[1] Being fermionic, G inherits the spin-structure question just settled: its moding r is half-integer in the NS sector and integer in the R sector—the same two choices, now inside the symmetry algebra itself. What the enlarged algebra buys (a square root of the Hamiltonian, and with it the sturdiest partition functions in these notes) is the closing business of Lecture 5.

The thread converges on SCFT

The choice "which spin structure" is the choice "insert $(-1)^F$ on which cycle." Inserting it on the time cycle of the R sector produces the trace $\mathrm{Tr}_R(-1)^F q^{L_0-c/24}$—the **Witten index** (Lecture 5). The theta functions we just met are precisely the super-characters that index will assemble. The boson gave us η; the fermion gives us the θ's; together they are the building blocks of superconformal partition functions.

4.5 Characters: the partition function decomposes

The Hilbert space organizes into Virasoro representations V_h labelled by conformal weight h. The trace over a single representation is its **character**,

$$\chi_h(\tau) = \mathrm{Tr}_{V_h}\, q^{L_0-c/24} = q^{h-c/24}\sum_{n\geq 0}(\dim V_{h,n})\, q^n, \tag{4.12}$$

and the full partition function is a sesquilinear combination

$$Z(\tau,\bar\tau) = \sum_{h,\bar h} N_{h\bar h}\,\chi_h(\tau)\,\overline{\chi_{\bar h}(\tau)}, \qquad N_{h\bar h}\in\mathbb{Z}_{\geq 0}. \tag{4.13}$$

[1] Indeed the pattern of §4.1 persists one rung up. Adjoin to z an anticommuting partner θ (with $\theta^2=0$), making the worldsheet a *superspace* with coordinates $(z\,|\,\theta)$; a function $f(z,\theta)=f_0(z)+\theta f_1(z)$ packs a boson and a fermion into one *superfield*. Superconformal transformations are exactly the *superholomorphic* maps of $(z\,|\,\theta)$—those preserving, up to scale, the odd derivative $D=\partial_\theta+\theta\,\partial_z$, a square root of translations ($D^2=\partial_z$) foreshadowing the square root of the Hamiltonian in Lecture 5. Mode-expanding one such map yields the L_m and G_r together: super-Virasoro is to superholomorphy what Virasoro is to holomorphy. We will not need superspace, but the hint is worth having.

In a rational theory there are finitely many χ_h, and the modular S-transform acts on this finite vector by a matrix—the **modular S-matrix** $S_{hh'}$. Modular invariance becomes a finite linear-algebra condition on the multiplicities $N_{h\bar{h}}$. Equation (4.12) is the doorway to Lecture 5: a character is manifestly a *graded dimension*.

4.6 Holomorphic factorization, and its instructive failure

Stare at the shape of (4.13). Every character $\chi_h(\tau)$ is *holomorphic* in τ: built from one chirality alone, it sees q and never $\bar{q}$, exactly as promised in §4.1. The partition function is therefore a finite sesquilinear form in holomorphic objects—chiral blocks and their conjugates, glued by a matrix of non-negative integers. Now recall the first structural fact of these notes (§1.4): *independent systems multiply.* If the left- and right-moving halves of the theory were independent systems, Z would be a single product, $Z(\tau, \bar{\tau}) = f(\tau)\,\overline{f(\tau)}$. Equation (4.13) says that in general they are not: no such f exists, and the matrix $N_{h\bar{h}}$ records exactly how the two chiral halves are correlated—which left-moving sectors occur, and paired with which right-moving ones. In information-theoretic terms N is a joint distribution over sector labels, and the failure of Z to factor is the statement that the chiral halves carry mutual information. The partition function has acquired *internal structure*: it is not one function but a pairing of two vectors of blocks.

Both running examples exhibit the failure. For the boson, (4.6) glues η^{-1} to $\bar{\eta}^{-1}$ diagonally, and the zero-mode factor $1/\sqrt{\operatorname{Im}\tau}$—neither holomorphic nor antiholomorphic—ties the halves together. For the fermion, the invariant is a *sum* over spin structures, $\sum_i |\theta_i/\eta|$, correlating the left and right boundary conditions sector by sector. Modular invariance is the culprit: S and T shuffle the blocks among themselves, so (special theories aside) no single product $f\bar{f}$ is invariant—only matrix-glued combinations survive. The diagonal gluing $N_{h\bar{h}} = \delta_{h\bar{h}}$ always works, because the S-matrix is unitary; but it is typically not alone, and distinct gluing matrices are *inequivalent theories built from identical chiral ingredients.* The chiral data does not determine the CFT: the pairing is extra physical information.

This is precisely why Lecture 5 studies the chiral half by itself. The blocks are characters of the *chiral algebra*—the affine Lie algebras and vertex operator algebras of the next lecture—and representation theory speaks about one chirality at a time. A full CFT is two chiral halves plus a modular-invariant gluing; the objects of Lecture 5 are the factors that (4.13) almost, but never quite, separates into. (For the reader keeping categorical score: the blocks span a finite-dimensional representation of $SL(2, \mathbb{Z})$, and a consistent Z is an *invariant vector* in a product of two such—a first hint of the finale's habit of valuing partition functions in vector spaces rather than numbers.)

For the curious: ADE, and when factorization does hold

For $\widehat{\mathfrak{su}}(2)_k$—and, closely related, for the Virasoro minimal models—the complete list of gluing matrices is known, and it follows a pattern no one ordered: the modular invariants are classified by the simply-laced Dynkin diagrams A_n, D_n, E_6, E_7, E_8 (Cappelli–Itzykson–Zuber). The A-series is the diagonal theory; the D-series pairs sectors with their images under a $\mathbb{Z}_2$ symmetry; among the exceptionals, E_6 and E_8 are diagonal invariants of an *extended* chiral algebra in disguise, while E_7 genuinely couples different left and right sectors. Honest holomorphic factorization—or an outright holomorphic Z—is reserved for *meromorphic* theories and requires $c \in 24\mathbb{Z}$, else phases like η's $\mathrm{e}^{\mathrm{i}\pi/12}$ survive under T. The most famous example sits at $c = 24$: the Moonshine module, with $Z(\tau) = J(\tau)$ the modular j-function (less its constant term) and the Monster group as its symmetry.

Rosetta notes

- $\tau \leftrightarrow$ complexified inverse temperature/geometry; $q^{L_0-c/24} \leftrightarrow \mathrm{e}^{-\beta H}$ with vacuum subtraction.
- Spin structure $\leftrightarrow$ L3 boundary conditions $\leftrightarrow$ $(-1)^F$ insertions.
- Modular $S \leftrightarrow$ high/low-temperature duality $\leftrightarrow$ the Kac–Peterson transform (L5). The elliptic variable y *is* the chemical potential of §2.4, now a Jacobi-form variable.
- Chiral block/character $\leftrightarrow$ a module of the chiral algebra (L5); the gluing matrix $N_{h\bar{h}} \leftrightarrow$ which left and right modules pair—the part of the theory the chiral halves alone cannot see.

Chapter 5

Representation Theory: partition functions as characters

For the reader from physics

"Module" is the state space of a sector; "character" is its partition function; "level k" is tied to your central charge. The Weyl–Kac character formula is the rep-theoretic ancestor of the CFT characters you just met, and the statement that characters transform into each other under $SL(2,\mathbb{Z})$ is, on the physics side, modular invariance—now a theorem.

Lecture 4 closed with a diagnosis (§4.6): the torus partition function is a matrix-glued pairing of *holomorphic blocks*, and the blocks—not the full Z—are where representation theory lives. This lecture is about the blocks. Everything in it concerns a single chirality at a time, and the reward for giving up the other half is that the objects become algebraic and the statements become theorems.

5.1 A character is a graded dimension is a partition function

Let $V = \bigoplus_{n\geq 0} V_n$ be a graded vector space with finite-dimensional pieces. Its **(graded) character**, or q-character, is

$$\mathrm{ch}_q V = \sum_{n\geq 0} (\dim V_n)\, q^n. \tag{5.1}$$

Here q is a *formal* variable—a bookkeeping device whose sole job is to file the integer $\dim V_n$ at position n. Identities between characters are identities of formal power series; no convergence is required or asked. Physics re-enters by *specialization*: set $q = \mathrm{e}^{2\pi \mathrm{i}\tau}$ (the nome of Lecture 4) and the bookkeeping becomes Boltzmann weighting, the series now converging for $|q| < 1$. This is the same object as the CFT character (4.12) (up to the $q^{h-c/24}$ prefactor) and—reading Lecture 2 backwards—the same object as a partition function: $\dim V_n$ counts the states at grade n, and q^n is their Boltzmann weight.

One signpost before the machinery: how precisely do ch and the physical Z relate? First, keep the two symbols apart. $\mathrm{ch}\, V$ is the character of a single module V: a *formal* series, possibly in many variables e^{λ}. The block χ of Lecture 4 is the same object specialized to $q = \mathrm{e}^{2\pi \mathrm{i}\tau}$ and dressed with the modular anomaly,

$$\chi_\Lambda(\tau) \;=\; q^{\,h_\Lambda - c/24}\, \mathrm{ch}_q\, L(\Lambda),$$

where h_Λ is the ground-state weight of the sector (*Lecture* 4 labelled blocks by h; here the label is the highest weight Λ, which determines h_Λ). So ch is the algebraic animal, χ its modular avatar, and each is a *chiral* partition function: the trace over one sector, at one chirality. The whole construction of this lecture, as one cartoon:

$$\underbrace{(\mathfrak{g},\,k)}_{\text{chiral data}} \rightsquigarrow \underbrace{\{L(\Lambda)\}}_{\text{sectors}} \xrightarrow{\text{ch},\ q^{h_\Lambda - c/24}} \underbrace{\chi_\Lambda(\tau)}_{\text{blocks}} \xrightarrow{\text{glue}} Z(\tau,\bar\tau) = \sum_{\Lambda,\bar\Lambda} N_{\Lambda\bar\Lambda}\,\chi_\Lambda(\tau)\,\overline{\chi_{\bar\Lambda}(\tau)}. \tag{5.2}$$

Defining the full theory means *summing over all the available representations*—the sector list. And the integer matrix $N_{\Lambda\bar\Lambda}$ is no extra axiom but the theory's *spectrum*, presented as a matrix: the Hilbert space decomposes under the two chiral algebras as $\mathcal{H} = \bigoplus_{\Lambda,\bar\Lambda} N_{\Lambda\bar\Lambda}\, L(\Lambda)\otimes\overline{L(\bar\Lambda)}$, with $N_{\Lambda\bar\Lambda}$ counting how many times the sector $(\Lambda,\bar\Lambda)$ occurs; the graded trace over this $\mathcal{H}$ produces (5.2) verbatim. Modular invariance (§4.6) is then the consistency condition selecting which matrices N define theories at all. The comparison with Lecture 3 runs deeper than bookkeeping: it says *what specifies a model*. There, a theory is specified by field content and an action, and Z is computed from that data, $S_E \rightsquigarrow Z = \int \mathcal{D}\phi\, \mathrm{e}^{-S_E}$; here, by a chiral algebra and a gluing, $(\mathfrak{g}_k,\,N) \rightsquigarrow Z$ via (5.2). The algebra plays the role the action played: the input that defines the model. For the WZW models the two specifications literally coincide—the same theory may be handed to you as a sigma-model action on a group manifold or as $\widehat{\mathfrak{g}}$ at level k, and the partition function does not care which. The rest of the lecture constructs the ingredients of (5.2): the next two sections build $L(\Lambda)$ and ch, the modularity section explains why the blocks mix under $SL(2,\mathbb{Z})$, and Example 5.2 assembles a complete Z.

The finite-dimensional Weyl character formula is the classical ancestor; we need its affine extension.

5.2 A Lie-theory refresher, on the lattice

The Weyl–Kac formula below is the centrepiece of this lecture, and on first sight it is a dense block of undefined symbols. Here is the minimum needed to read it—worth skimming even if Lie theory is home turf, since it fixes notation. Keep $\mathfrak{g} = \mathfrak{su}(2)$ in mind throughout.

The Cartan subalgebra is the set of good quantum numbers. A **simple** Lie algebra $\mathfrak{g}$ is a finite-dimensional *non-abelian* Lie algebra with no nontrivial ideals—nothing invariant to quotient away ($\mathfrak{su}(2)$, $\mathfrak{su}(n)$, …). Inside it choose a **Cartan subalgebra** $\mathfrak{h}$: a maximal set of commuting, simultaneously diagonalizable generators—physically, a maximal set of compatible conserved charges. For $\mathfrak{su}(2)$ it is the span of J^3; the dimension of $\mathfrak{h}$ is the **rank**.

Roots, weights, and their lattices. Now diagonalize everything against $\mathfrak{h}$. The algebra acting on itself sorts into eigenvectors: the nonzero eigenvalues $\alpha \in \mathfrak{h}^*$ are the **roots**, and the root vectors are ladder operators (the $J^\pm$ of $\mathfrak{su}(2)$). A representation likewise sorts into simultaneous eigenstates, and the eigenvalue $\lambda \in \mathfrak{h}^*$ of a state is its **weight**—nothing but its tuple of quantum numbers (for $\mathfrak{su}(2)$, the magnetic number m). The weights of finite-dimensional representations populate the **weight lattice** P; the roots generate the **root lattice** $Q \subseteq P$. A finite-dimensional irreducible is labelled by its **highest weight** Λ (the spin j), and for character purposes a representation is nothing more than its list of weights, counted with repetition:

$$\operatorname{ch} V_\Lambda = \sum_\lambda \operatorname{mult}(\lambda)\,\mathrm{e}^\lambda,$$

where $\operatorname{mult}(\lambda)$ is the number of independent states carrying the quantum numbers λ—one formal exponential per state, a Boltzmann weight per conserved charge, as the box below insists. Mind the two roles here: the weight λ is a *label*, naming which quantum numbers a state carries, while the *variable* is a Cartan element $h \in \mathfrak{h}$, against which each formal symbol evaluates to a number, $\mathrm{e}^\lambda \mapsto \mathrm{e}^{\lambda(h)}$. For $\mathfrak{su}(2)$: $\operatorname{ch} V_j = \sum_{m=-j}^{j} z^m$, with $z = \mathrm{e}^\theta$ the fugacity conjugate to J^3—the weights m label the terms; z is the variable. For the irreducibles of $\mathfrak{su}(2)$ every multiplicity is 1: spin

j has weights $m = -j, \dots, +j$, each occurring once. Multiplicity starts carrying information the moment representations combine or the rank grows. In $\frac{1}{2} \otimes \frac{1}{2}$ the weight $m = 0$ occurs twice—$\mathrm{mult}(0) = 2$, one state from the triplet and one from the singlet. In the adjoint of $\mathfrak{su}(3)$ (the octet), the six nonzero weights are the roots, each with multiplicity 1, while the weight 0 has $\mathrm{mult}(0) = 2$: the two Cartan directions—in the meson octet, the π^0 and the η.

The Weyl group is a point group. The **Weyl group** W is the finite group generated by reflections in the hyperplanes orthogonal to the roots; it permutes the weights of any representation. For $\mathfrak{su}(2)$ it is $m \mapsto -m$, so $W = \mathbb{Z}_2$. For $\mathfrak{su}(3)$ (rank 2) the six roots form a regular hexagon and $W = S_3$: the six-element group permuting the three eigenvalues of the fundamental, generated by the two reflections s_1, s_2 in the walls. The **length** $\ell(w)$ is the minimal number of generating reflections composing w, and $(-1)^{\ell(w)}$ is its sign—concretely, for $\mathfrak{su}(3)$:

$w \in W(\mathfrak{su}(3)) = S_3$	geometrically	$\ell(w)$	$(-1)^{\ell(w)}$
e	identity	0	$+$
$s_1, \quad s_2$	the two wall reflections	1	$-$
$s_1 s_2, \quad s_2 s_1$	rotations by $\pm 120°$	2	$+$
$s_1 s_2 s_1 = s_2 s_1 s_2$	the remaining reflection	3	$-$

so the sign is nothing but $\det w$: reflections count -1, rotations $+1$. Finally the **Weyl vector** ρ, half the sum of the positive roots, produces the ubiquitous quantum shifts (for $\mathfrak{su}(2)$, $j \to j + \frac{1}{2}$). A condensed-matter reader has seen this entire package before: a lattice and its *point group*. What is missing is translations—and translations are exactly what the affine extension supplies.

The affine extension is Fourier analysis on the circle of Lecture 4. Put the theory on the space circle. A $\mathfrak{g}$-valued symmetry current becomes a $\mathfrak{g}$-valued function on the circle, and its Fourier modes $X \otimes t^m$ (with $t = \mathrm{e}^{\mathrm{i}\sigma}$) span the **loop algebra** $\mathfrak{g} \otimes \mathbb{C}[t, t^{-1}]$. Quantum mechanically the modes acquire a central term—the 2d sibling of a Schwinger term:

$$[\, X \otimes t^m,\ Y \otimes t^n \,] \;=\; [X, Y] \otimes t^{m+n} \;+\; m\, \langle X, Y \rangle\, \delta_{m+n,0}\, K, \tag{5.3}$$

with $\langle \cdot, \cdot \rangle$ the suitably normalized Killing form and K central. This is the **affine Lie algebra** $\widehat{\mathfrak{g}}$. On an irreducible highest-weight module K acts by a scalar k: the **level**. (Nothing yet restricts k; the integrability of the next section is what forces $k \in \mathbb{Z}_{\geq 0}$.) The Cartan grows to $\mathfrak{h} \oplus \mathbb{C}K \oplus \mathbb{C}d$, where d grades by Fourier mode number—L_0 in Lie clothing—so an affine weight is a triple (finite weight, level, energy grade), and the character variable q of (5.1) is the fugacity conjugate to d.

The affine Weyl group is a space group. The reflections compatible with the extension assemble into $W_{\mathrm{aff}} = W \ltimes Q^\vee$: the finite Weyl group, semidirect *translations* by the (co)root lattice. To the condensed-matter eye this is precisely the crystallographer's upgrade from point group to **space group**. The translations are also why theta functions are imminent: a sum of exponentials of a quadratic form over a lattice *is* a theta function.

5.3 Integrable modules and the Weyl–Kac formula

Which representations of $\widehat{\mathfrak{g}}$ play the role the finite-dimensional irreducibles played for $\mathfrak{g}$? The **integrable** highest-weight modules $L(\Lambda)$: those on which every root's $\mathfrak{su}(2)$ exponentiates to the group (equivalently, Λ is dominant integral)—the affine analogue of finite-dimensionality. Integrability is a strong finiteness condition: at fixed level k there are only finitely many such Λ—for $\widehat{\mathfrak{su}}(2)_k$, the spins $j = 0, \frac{1}{2}, \dots, \frac{k}{2}$—and they are exactly the sectors of the level-k WZW model.

In field-theory language, $L(\Lambda)$ is the chiral-algebra avatar of *a quantum field transforming in the representation* Λ *of* $\mathfrak{g}$: its ground states form the finite-dimensional $\mathfrak{g}$-module Λ, created by the corresponding primary field, and the rest of the module is that field's cloud of excitations—the current descendants J^a_{-n} acting on the ground states. One module = one field, together with everything the symmetry algebra can do to it. Their characters are given by the **Weyl–Kac character formula**,

$$\operatorname{ch} L(\Lambda) = \frac{\sum_{w\in W}(-1)^{\ell(w)}\,\mathrm{e}^{w(\Lambda+\rho)-\rho}}{\prod_{\alpha>0}(1-\mathrm{e}^{-\alpha})^{\operatorname{mult}\alpha}}, \tag{5.4}$$

the affine analogue of Weyl's finite-dimensional formula—and every ingredient is now in hand: W is the affine Weyl group of §5.2, $\ell(w)$ its length, ρ its Weyl vector, and $\operatorname{mult}\alpha$ the dimension of the root space (1 for the "real" roots—the finite roots and their translates up the energy grading; $\operatorname{rank}\mathfrak{g}$ for the "imaginary" roots, the pure energy directions). Because $W_{\text{aff}} = W \ltimes Q^\vee$, the numerator reorganizes as a finite sum over W of *lattice sums*: specializing the formal exponentials to the single grading variable q, the character becomes a finite sum of *string functions* times *theta functions*—the theta functions of Lecture 4 reappearing as honest characters, their lattice-sum forms in (4.8) now explained rather than merely checked.

The Cartan variables are chemical potentials

The character (5.4) is a function *on the Cartan subalgebra*: evaluated at $h \in \mathfrak{h}$, each formal exponential becomes the number $\mathrm{e}^{\lambda(h)}$—the weights λ label the terms, the components of h are the variables, rank-many of them plus q, exactly as in §5.2. Those variables are precisely the chemical potentials/fugacities of §2.4: each Cartan direction is a conserved charge, and its conjugate variable grades the character. The principal specialization (setting the Cartan variables to special values, keeping only $q = \mathrm{e}^{2\pi\mathrm{i}\tau}$ conjugate to the energy L_0) is "turning the chemical potentials off" to leave the pure q-grading of (5.1).

5.4 Modularity, as a theorem

The miracle—discovered by physicists as modular invariance and proved by mathematicians as a property of characters—is:

> **Kac–Peterson.** The normalized characters of the level-k integrable modules of $\widehat{\mathfrak{g}}$ span a finite-dimensional representation of $SL(2,\mathbb{Z})$. Under $S : \tau \mapsto -1/\tau$ they transform into one another by a unitary matrix $S_{\Lambda\Lambda'}$, and under $T : \tau \mapsto \tau + 1$ by a diagonal phase.

The $-c/24$ shift (with $c = k\dim\mathfrak{g}/(k+h^\vee)$ the Sugawara central charge) is *exactly* the correction that makes the characters transform as vector-valued modular forms. The physicist's modular S-matrix of §4.5 *is* the Kac–Peterson matrix. Thread 3 closes: the regularized vacuum energy of Lecture 3 is the modular anomaly of Lecture 5.

Example 5.1 (Kac–Peterson for $\widehat{\mathfrak{su}}(2)_k$, explicitly)**.** With sectors labelled by $j, j' \in \{0, \frac{1}{2}, \dots, \frac{k}{2}\}$,

$$S_{jj'} = \sqrt{\frac{2}{k+2}}\,\sin\Bigl(\frac{(2j+1)(2j'+1)\,\pi}{k+2}\Bigr), \qquad T_{jj'} = \delta_{jj'}\,\mathrm{e}^{2\pi\mathrm{i}(h_j - c/24)}, \quad h_j = \frac{j(j+1)}{k+2}. \tag{5.5}$$

At level 1 (two sectors; $h_0 = 0$, $h_{1/2} = \frac{1}{4}$, $c = 1$):

$$S = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix}, \qquad T = \begin{pmatrix} \mathrm{e}^{-\pi\mathrm{i}/12} & 0 \\ 0 & \mathrm{e}^{5\pi\mathrm{i}/12} \end{pmatrix}. \tag{5.6}$$

S is real, symmetric, unitary, and squares to $\mathbf{1}$. Modular invariance of the diagonal theory $\sum_j |\chi_j|^2$ (assembled below in Example 5.2) follows at a glance: under S the sum is preserved because S is unitary; under T the phases cancel between χ_j and $\overline{\chi_j}$. The diagonal invariant is modular invariant *because* Kac–Peterson holds.

5.5 Vertex operator algebras: making Lecture 4 rigorous

A **vertex operator algebra** (VOA) V axiomatizes the chiral algebra of a 2d CFT: a graded space with a state–field correspondence $a \mapsto Y(a,z)$, a vacuum, and a *conformal vector* ω whose field is the stress tensor,

$$T(z) \;=\; Y(\omega, z) \;=\; \sum_n L_n\, z^{-n-2},$$

its modes realizing the Virasoro algebra (4.1) on the state space. The connection between vector and symmetry is tight in both directions. As a state, $\omega = L_{-2}|0\rangle$ has weight 2 and carries the central charge on board: $L_2\,\omega = \frac{c}{2}|0\rangle$. As a field, the commutators (4.1) are equivalent to the singular part of the **operator product expansion** (OPE) of the stress tensor with itself,

$$T(z)\,T(w) \;\sim\; \frac{c/2}{(z-w)^4} + \frac{2\,T(w)}{(z-w)^2} + \frac{\partial T(w)}{z-w},$$

the same algebra said pointwise. This is how a VOA *manifests* conformal symmetry: not as transformations of a spacetime, but as a distinguished vector whose modes act on everything. Its modules are the sectors, and the graded dimension of a module is the chiral torus partition function (4.12). The rigorous form of modular invariance is:

> **Zhu's theorem.** For a sufficiently nice VOA (rational and C_2-cofinite),[1] the span of the characters of its irreducible modules is invariant under $SL(2,\mathbb{Z})$.

This is the theorem standing behind the physical assertion of §4.2 that $Z(\tau)$ is modular invariant.

The running examples, identified

The **free boson** is the Heisenberg (Fock) VOA; its character $\mathrm{Tr}\, q^{L_0 - c/24}$ *is* $\eta(\tau)^{-1}$—the $\prod(1-q^n)^{-1}$ of Lecture 2 dressed in its regularized zero-point factor $q^{-1/24}$, now read as a module character. The **free fermion** is a Clifford-module (super) VOA; its NS and R modules are exactly the spin-structure sectors of Lecture 4, with characters the θ/η blocks (4.9). The single modes of Lecture 1 have become the generators of these vertex algebras. The loop is closed.

Example 5.2 (Three characters, expanded)**.** The generalities are best absorbed through leading terms; for each of three modules we list its modular block $\chi = q^{\,h-c/24}\,\mathrm{ch}_q$.

- The *Heisenberg VOA* $V_{\mathfrak{h}}$ (the free boson, $c=1$; the name honours its underlying one-dimensional abelian algebra $\mathfrak{h}$, of which it is the affinization):

$$\chi_{V_{\mathfrak{h}}} = q^{-1/24} \prod_{n\geq 1} \frac{1}{1-q^n} = q^{-1/24}(1 + q + 2q^2 + 3q^3 + 5q^4 + \cdots),$$

  the partition numbers of (2.7): one state per way of exciting the modes $a_{-1}, a_{-2}, \ldots$ of Lecture 2. Its conformal vector is $\omega = \frac{1}{2}\,a_{-1}^2|0\rangle$, whose field is $T = \frac{1}{2}\,{:}\partial\phi\,\partial\phi{:}$.

[1] *Rational*: finitely many irreducible modules, and every module a direct sum of them. *C_2-cofinite*: the subspace spanned by the vectors $a_{-2}b$ (for $a, b \in V$, with a_{-2} a Laurent mode of the vertex operator of a) has finite codimension in V—a finiteness condition that makes the characters converge and satisfy a modular differential equation. Together these are the standing hypotheses of rational VOA theory.

- The *Clifford (super) VOA* F (the free fermion, $c = \frac{1}{2}$), in its NS module $F_{\rm NS}$:

$$\chi_{F_{\rm NS}} = q^{-1/48} \prod_{n\geq 1} (1 + q^{n-\frac{1}{2}}) = q^{-1/48}(1 + q^{\frac{1}{2}} + q^{\frac{3}{2}} + q^2 + q^{\frac{5}{2}} + q^3 + q^{\frac{7}{2}} + 2q^4 + \cdots),$$

  partitions into *distinct* half-odd-integers, each mode ψ_{-r} usable once (Pauli again); note that grade 1 is unreachable. This is the block $\sqrt{\theta_3/\eta}$ of (4.9). Here $\omega = \frac{1}{2}\,\psi_{-3/2}\psi_{-1/2}|0\rangle$, whose field is $T = -\frac{1}{2} :\!\psi\,\partial\psi\!:$.

- $\widehat{\mathfrak{su}}(2)_1$*, vacuum module* (the simplest WZW chiral algebra; $c = \frac{k\dim\mathfrak{g}}{k+h^\vee} = \frac{3}{1+2} = 1$):

$$\chi_{j=0} = q^{-1/24}\, \frac{\sum_{n\in\mathbb{Z}} q^{n^2}}{\prod_{n\geq 1}(1-q^n)} = q^{-1/24}(1 + 3q + 4q^2 + 7q^3 + 13q^4 + \cdots),$$

  a theta function on the $\mathfrak{su}(2)$ root lattice over the oscillator product—the string-function-times-theta shape of the Weyl–Kac formula, in the flesh. Read the coefficients as $\mathfrak{su}(2)$ representations: the 3 at grade one is the adjoint, the current states $J^a_{-1}|0\rangle$ themselves; the 4 at grade two is $\mathbf{3} \oplus \mathbf{1}$—the spin-2 combination of two currents is a *null state* at level 1, integrability visibly pruning the spectrum. Restoring a Cartan fugacity $z^{J^3_0}$ would resolve each coefficient into its weights: the Cartan-variables box, in action. And lest the vacuum module float free of its theory: level 1 has exactly two integrable spins, and the full WZW torus partition function of Lecture 4 glues them diagonally,

$$Z_{\widehat{\mathfrak{su}}(2)_1} = \sum_{j\,=\,0,\,\frac{1}{2}} \chi_j\,\overline{\chi_j},$$

  the sesquilinear form (4.13) with $N = \mathbf{1}$—one full theory, assembled from its two chiral blocks: the cartoon (5.2), realized end to end. The conformal vector here is nobody new, just the currents squared: the **Sugawara** construction $\omega = \frac{1}{2(k+h^\vee)}\sum_a J^a_{-1}J^a_{-1}|0\rangle = \frac{1}{6}\sum_a \left(J^a_{-1}\right)^2|0\rangle$—the origin of the central-charge formula $c = k\dim\mathfrak{g}/(k+h^\vee)$ quoted above.

Marrying the first two examples pays a dividend: $V_{\mathfrak{h}} \otimes F$ carries, besides $\omega = \omega_{\rm b} + \omega_{\rm f}$, the *superconformal* vector $G = a_{-1}\psi_{-1/2}|0\rangle$, whose field is exactly the supercurrent $G = \psi\,\partial\phi$ of Lecture 4. That super-VOA is where §5.6 begins.

5.6 Contact with superconformal field theory

What is the sturdiest partition function of all? This closing section builds supersymmetry into the chiral algebra in two steps, and each step buys a more protected object: first an integer, the **Witten index**; then a Jacobi form, the **elliptic genus**.

Step one is to manifest, in VOA language, the superconformal symmetry met at the close of Lecture 4 (§4.4)—and it costs nothing we do not already own. Just as the conformal vector ω of §5.5 realizes Virasoro, a **super-VOA** distinguishes one more vector. In a theory with the right matter—the boson-plus-fermion system at $c = \frac{3}{2}$—the state $G = \psi\,\partial\phi$ already sits in the weight-$\frac{3}{2}$ graded piece, counted by the partition function all along, just as the affine currents J^a of §5.2 sit at weight 1. Promote it, from a state one merely counts to a generator one organizes by: the state–field correspondence returns exactly the supercurrent $G(z)$ of Lecture 4, its modes G_r closing with the L_m on the super-Virasoro algebra (4.11). In the OPE language of §5.5 the package reads: T sees G as a primary of weight $\frac{3}{2}$, and G multiplies with itself back into the stress tensor,

$$G(z)\,G(w) \;\sim\; \frac{2c/3}{(z-w)^3} + \frac{2\,T(w)}{z-w},$$

the pointwise form of (4.11)'s anticommutator—and the local seed of the square root of the Hamiltonian, coming next. Because G is fermionic, the dichotomy of §3.2 applies to it: its moding runs over half-integers (NS) or integers (R)—the two spin structures, now internal to the chiral algebra—and the modules of the super-VOA come in NS and R families accordingly.

Promotion changes the bookkeeping, not the theory

Promoting a current adds nothing: the states and Z are untouched. What changes is the *decomposition*—the same spectrum reassembles into fewer, larger modules, and the cartoon (5.2) re-runs with a bigger algebra on the left. The rhyme with §5.2 is exact: promote weight-1 currents and get $\widehat{\mathfrak{g}}$; promote a weight-$\frac{3}{2}$ fermionic current and get super-Virasoro. A chiral algebra is specified by which of the theory's currents one takes as generators—the "action" of §5.1, once more.

The two sectors are not symmetric: only the R sector, with its integer moding, contains a zero mode G_0, and setting $r = s = 0$ in the anticommutator of (4.11) gives

$$G_0^2 \;=\; L_0 - \frac{c}{24}. \tag{5.7}$$

The R sector owns a *square root of the chiral Hamiltonian*—built on the very $-c/24$ of our thread. Two consequences follow at once. First, the chiral energy $L_0 - \frac{c}{24}$ is non-negative on a unitary R module: it is a square. Second, every state of *positive* chiral energy is paired: G_0 maps it to a partner of equal energy and opposite fermion parity, and since applying G_0 twice returns a nonzero multiple of the original state, neither partner can vanish. Only zero-energy states—those annihilated by G_0—can sit alone.

Now take the sign-graded trace of Lecture 2 over an R module—the **Witten index**

$$\mathrm{Tr}_R\,(-1)^F\,q^{\,L_0-c/24}. \tag{5.8}$$

At every positive energy the paired states contribute with opposite signs and cancel, so the entire q-series collapses to a single integer, independent of β (equivalently of $\mathrm{Im}\,\tau$): the number of zero-energy ground states, counted with sign. This is exactly the "net count of unpaired ground states" promised by the vacuum-energy cancellation of Lecture 3, and it is rigid: under any continuous deformation preserving the algebra, states can leave or reach zero energy only in G_0-pairs of opposite parity, so the signed count cannot move. A partition function has become a *deformation invariant.*

Geometrically we have met this object already. Inserting $(-1)^F$ makes the *time* circle periodic (§3.2), so (5.8) is the (R,R) torus amplitude—the fourth spin structure of Lecture 4, the one that evaluated to $\sqrt{\theta_1/\eta} \equiv 0$. The free Majorana fermion shows why in miniature: its two R ground states carry opposite fermion parity, so they cancel and the index vanishes. $\theta_1 \equiv 0$, recomputed by counting.

A vanishing index is protected but not talkative. To refine it, promote once more. Suppose the chiral algebra can be enlarged by a weight-1 $U(1)$ current J under which the supercurrent splits into oppositely charged halves $G^{\pm}$: the $\mathcal{N}=2$ **super-Virasoro algebra**, its $U(1)$ the *R-symmetry.* (Free fields manage this with complex matter—a complex boson paired with a complex fermion, $c = 3$.) The zero mode J_0 is one more conserved charge, so the Interlude's reflex fires: hand it a conjugate potential, the fugacity y. The refined trace is the **elliptic genus**

$$\boxed{\mathcal{E}(\tau,y) = \mathrm{Tr}_R\,(-1)^F\,y^{J_0}\,q^{\,L_0-c/24},} \tag{5.9}$$

and the trace need no longer collapse: states that cancelled in (5.8) by parity can carry different J_0 charges, so they now cancel only at $y = 1$—the fugacity *revives* what $\theta_1 \equiv 0$ set to zero. What deformation now protects is a whole function, a weak Jacobi form in (τ, y) once the $U(1)$ charges are suitably integral—for a sigma model, a Calabi–Yau target M. The name promises geometry, and delivers: $\mathcal{E}$ is then a genuine *invariant* of the smooth compact manifold M. Like the Witten index it is protected against deformation, and it bears real information: special values of y recover the Euler characteristic, the Hirzebruch signature and the $\hat{A}$-genus of M, which the q-expansion binds into a single modular object. A partition function has become a manifold invariant—index theory and modular forms shaking hands. Every thread of these notes meets here: the graded trace (L2), the sign grading / spin structure (L3–L4), the modular transformation (L4–L5), and the chemical potential y (the Interlude) all fuse into a single modular object, and the theta functions of Lecture 4 reappear as its super-characters.

For the curious: which supersymmetry, and a Euclidean subtlety

The supersymmetry of this section is *worldsheet* supersymmetry—the two-dimensional super-Virasoro algebra, pairing a boson with a fermion on the worldsheet (the mode-by-mode cancellation of Lecture 3)—not the spacetime super-Poincaré algebra of particle physics: the index and the genus are statements about a 2d chiral algebra, whatever spacetime physics the model describes. The Euclidean framing adds its own subtlety. Where Lorentzian supersymmetry leans on the reality $\{Q, Q^\dagger\} = H \geq 0$, the Euclidean theory lets the holomorphic and antiholomorphic supercharges go their separate ways, and it is that holomorphic structure—not positivity—that makes the index and the genus *modular* objects: Jacobi forms, not merely deformation-invariant numbers.

Rosetta notes

- Character $\leftrightarrow$ partition function $\leftrightarrow$ graded dimension.
- Level k $\leftrightarrow$ central charge c; Kac–Peterson S $\leftrightarrow$ CFT modular S $\leftrightarrow$ high/low-temperature duality.
- R-sector $(-1)^F$ trace $\leftrightarrow$ the Witten index $\leftrightarrow$ a supersymmetric partition function $\leftrightarrow$ the elliptic genus.

Chapter 6

Finale: the partition function as a functor

For the reader from either camp

One sentence of category theory turns the translations you have been doing by hand into a single structural statement—and then, if you press it, admits that it has been describing only half of the series. The other half is one dimension up. This chapter proves a theorem, discovers what the theorem cannot reach, and goes to find it.

6.1 Functorial quantum field theory

Atiyah's axioms define a *topological* field theory in any dimension d: a functor out of Bord_d, whose objects are closed $(d-1)$-manifolds and whose morphisms are d-dimensional bordisms. We will work in $d = 2$ throughout—but not as a scoping convenience, and it is worth being exact about why, because the reason is the reason this chapter can prove a theorem instead of sketching a framework.

In dimension two the objects are *classified*. A closed 1-manifold is a disjoint union of circles and nothing else; there is one connected object, and the whole category is generated by a short list of surfaces glued along it. That is why Bord_2 is completely understood—in the topological case it is equivalent to the category of commutative Frobenius algebras, a theorem with a proof one can read in an afternoon—and it is why §6.2 below states a theorem rather than a research programme. Move up one dimension and the objects become surfaces, classified by genus, and the theory becomes rich enough to need a *category* of labels rather than a set; that is §6.4, and it is where the missing half of our story turns out to live. Move up two more and nobody has the objects under control at all. Dimension two is not where we are standing because the view is easy. It is where the ground is solid.

The second restriction is subtler and will matter in the same place. Atiyah's bordisms are bare manifolds; Segal's, in dimension two, carry a *conformal structure*, and the functor is then a conformal field theory rather than a topological one.[1] The distinction looks like bookkeeping and is not: it is exactly what permits the infinite-dimensional state space of Lecture 4, as the box at the end of §6.2 explains. Hold on to it. Everything in this chapter turns on which structure the bordisms are carrying.

With that said, a 2d (conformal) field theory *is* a symmetric monoidal functor

$$Z\colon\quad \mathrm{Bord}_2 \longrightarrow \mathrm{Vect}_{\mathbb{C}}\,. \tag{6.1}$$

[1]M. Atiyah, *Topological quantum field theories*, Publ. Math. IHÉS **68** (1988); G. Segal, *The definition of conformal field theory* (1988/2004). The two-dimensional topological case is proved in complete detail in J. Kock, *Frobenius Algebras and 2D Topological Quantum Field Theories*, LMS Student Texts **59**, Cambridge (2004).

First the vocabulary, precisely. The objects of Bord_2 are closed 1-manifolds—disjoint unions of circles—to be read as *space*. A morphism from Σ_1 to Σ_2 is a **bordism**: a compact surface M whose boundary is $\Sigma_1 \sqcup \Sigma_2$, to be read as a *spacetime* interpolating between them. (The literature says *cobordism* interchangeably; despite appearances, the "co" is a historical accident—"jointly bounding"—and not a categorical dual of anything.) The functor then assigns

- to a circle S^1, a Hilbert space $\mathcal{H}$ (the state space of Lecture 2);
- to a disjoint union of circles, the tensor product of their spaces (the factorization of Lecture 1);
- to a bordism, a linear map between those spaces (the path integral of Lecture 3 on that surface); and
- to the gluing of bordisms, the composition of the corresponding maps (the sewing of Lecture 3).

6.2 Why the partition function is a trace—now a theorem

Topologically, the torus T^2_τ of Lecture 4 is a cylinder with its two ends glued. As a bordism, the cylinder of modulus τ runs from one circle back to the same circle—a spacetime—and we know exactly which linear map the functor assigns it: the imaginary-time evolution of Lecture 2, dressed as the Boltzmann operator of Lecture 4,

$$Z(\mathrm{cylinder}_\tau) \;=\; q^{\,L_0 - c/24}\,\bar{q}^{\,\bar{L}_0 - c/24}.$$

Gluing the two ends closes the spacetime, and the functor Z sends "glue the ends of M" to "take the categorical **trace** of $Z(M)$." Hence

$$\mathsf{T}^2_\tau \;\longmapsto\; \mathrm{Tr}_{\mathcal{H}}\Big(q^{\,L_0 - c/24}\,\bar{q}^{\,\bar{L}_0 - c/24}\Big), \tag{6.2}$$

which is exactly (4.2). Said once and for all:

> **Theorem (gluing is trace).** Let $Z\colon \mathrm{Bord}_2 \to \mathrm{Vect}_{\mathbb{C}}$ be a symmetric monoidal functor, Σ an object, and M a bordism from Σ to Σ. Let $\widehat{M}$ denote the closed surface obtained by gluing the two ends of M to each other. Then
>
> $$Z(\widehat{M}) \;=\; \mathrm{Tr}\; Z(M).$$

The proof is a picture: inside Bord_2 the glued surface factors as $\widehat{M} = \mathrm{ev} \circ (M \sqcup \mathrm{id}) \circ \mathrm{coev}$—a composite of M with the two elbow half-cylinders of the box below—and a monoidal functor preserves the composite, which in $\mathrm{Vect}_{\mathbb{C}}$ *is* the trace.[2] The slogan of the entire series—*partition function = trace*—is therefore not a definition but a theorem: a consequence of functoriality alone. The trace of Lecture 2, the closed-surface path integral of Lecture 3, the torus amplitude of Lecture 4 are one statement, and we have just proved it.

And we have proved it about a *number*. Keep that in view for one more section, because it is about to become the problem.

For the curious: dualizability—the fine print promised in Lecture 2

When is a categorical trace defined at all? In a symmetric monoidal category, an object X is **dualizable** if it has a dual $X^\vee$ with evaluation $X^\vee \otimes X \to \mathbf{1}$ and coevaluation $\mathbf{1} \to X \otimes X^\vee$ obeying the zigzag identities; the trace of an endomorphism is built by composing the two. In Bord_2 *every* circle is dualizable: the elbow-shaped half-cylinders provide evaluation

[2] The gluing axiom is Atiyah's; the two-dimensional case is proved in complete detail in Kock, cited above.

and coevaluation, and the zigzag identity is literally "straighten the bent cylinder." Since a symmetric monoidal functor preserves duals, and the dualizable objects of $\mathrm{Vect}_{\mathbb{C}}$ are exactly the *finite-dimensional* spaces, a genuinely topological theory is forced to have finite-dimensional state spaces—and its $Z(\text{torus})$ needs no analysis whatsoever. The CFT of Lecture 4 evades this conclusion honestly: its bordisms carry conformal structure, every annulus has positive modulus—there is no infinitely thin elbow, and no identity cylinder among the bordisms—so $\mathcal{H}$ is never asked to be dualizable. Its infinite-dimensionality is paid for analytically instead: the cylinder map $q^{L_0-c/24}\,\bar{q}^{\bar{L}_0-c/24}$ is *trace class* for $|q|<1$, its trace converging courtesy of the Boltzmann suppression $\mathrm{e}^{-\beta E}$—the oldest object in these notes, doing the newest job.

6.3 Why the blocks are not functions

The theorem of §6.2 is true, and it accounts for less of these notes than its confidence suggests. Return to §4.6, where the torus partition function refused to factorize:

$$Z(\tau,\bar{\tau}) \;=\; \sum_{h,\bar{h}} N_{h\bar{h}}\,\chi_h(\tau)\,\overline{\chi_{\bar{h}}(\tau)}, \qquad N_{h\bar{h}}\in\mathbb{Z}_{\geq 0}.$$

We drew a moral there and then walked away from it. The chiral blocks χ_h are not functions that happen to sum to Z; they are the theory's actual content, and Z is a *bilinear pairing* of a vector of blocks with its conjugate, glued by a matrix that the blocks alone cannot predict. Lecture 5 was written to study that vector: its entries are characters of chiral-algebra modules, and Kac–Peterson says they span a finite-dimensional representation of $SL(2,\mathbb{Z})$. The number Z is an $SL(2,\mathbb{Z})$-*invariant* of that representation. The representation is the primary object; the invariant is a shadow it casts.

Now ask what the functor of §6.1 sees of any of this. Its value on the torus is a complex number, one per τ. It has no slot for a vector of blocks, no slot for an $SL(2,\mathbb{Z})$ representation, and—most tellingly—no way even to *ask* the question, because the mapping class group of the torus acts on the *object* T^2_τ of Bord_2, and the functor's value there is a scalar on which a group can act only trivially or not at all. Modular invariance, in this formalism, is a constraint we impose from outside. It is not something the category produces.

There is a second symptom, and the reader has been quietly tolerating it since §5.4. The characters do not transform under $SL(2,\mathbb{Z})$ honestly; they transform up to *phases*—the $\mathrm{e}^{-\pi\mathrm{i}/12}$ in the T-matrix of Example 5.1, the $\sqrt{-\mathrm{i}\tau}$ in $\eta(-1/\tau)$. A projective representation, not a linear one. Every formalism in these notes has swallowed those phases as a technicality. They are not a technicality. They are the fingerprint of a structure one level up, and they will be explained by it.

So: the blocks are vectors, they carry a projective $SL(2,\mathbb{Z})$ action, and $\mathrm{Bord}_2\to\mathrm{Vect}$ cannot house either fact. The object we want assigns a *vector space* to a torus—not a number. A functor assigns vector spaces to *objects*. Therefore the torus must be an object, not a closed morphism; therefore surfaces must be codimension one; therefore we are in three dimensions.

6.4 One dimension up: Chern–Simons and the block space

Take the bordism category of §6.1 and raise every dimension by one. The objects of Bord_3 are closed surfaces Σ—our tori, now promoted from spacetimes to spaces. The morphisms are 3-manifolds bounding them. A three-dimensional topological field theory is a symmetric monoidal functor

$$Z_3\colon\;\; \mathrm{Bord}_3 \longrightarrow \mathrm{Vect}_{\mathbb{C}}, \tag{6.3}$$

and it assigns to a surface a vector space,

$$\Sigma \;\longmapsto\; Z_3(\Sigma) \;=\; \mathcal{V}(\Sigma),$$

the **space of conformal blocks**. This is the vector §6.3 demanded, now produced rather than imposed. And because Bord_3 is genuinely topological—no conformal structure on the bordisms, unlike Lecture 4—the dualizability argument of the box above applies with full force and *bites*: $\mathcal{V}(\Sigma)$ is finite-dimensional. The same categorical machinery that let $\mathcal{H}$ be infinite-dimensional in two dimensions forbids it in three, and it is right both times. The block space is finite; the state space was not; the formalism knew the difference all along.

The theory in question is not abstract. It is **Chern–Simons theory** with gauge group G at level k, defined by the action

$$S_{\mathrm{CS}}[A] \;=\; \frac{k}{4\pi}\int_{M_3} \mathrm{Tr}\Big(A\wedge \mathrm{d}A + \tfrac{2}{3}\,A\wedge A\wedge A\Big), \tag{6.4}$$

whose Lagrangian is first order in derivatives and uses no metric—so the theory is topological by construction, and its Hilbert space on a surface Σ is obtained by quantizing the finite-dimensional moduli space of flat G-connections on Σ. The result, and the source of everything in this section, is Witten's:[3]

> **The blocks are a Hilbert space.** The Chern–Simons Hilbert space $\mathcal{V}_{G_k}(\Sigma)$ of a surface Σ is canonically the space of conformal blocks of the $\widehat{\mathfrak{g}}_k$ WZW model on Σ.

On the torus the statement is one we can check against Lecture 5 by counting. Quantizing flat connections on T^2 produces one state for each integrable highest weight of $\widehat{\mathfrak{g}}_k$—for $\widehat{\mathfrak{su}}(2)_k$, one for each spin $j = 0, \frac{1}{2}, \ldots, \frac{k}{2}$—so

$$\dim \mathcal{V}_{\widehat{\mathfrak{su}}(2)_k}(\mathsf{T}^2) \;=\; k+1,$$

and the distinguished basis of that space is precisely the list of blocks $\chi_j(\tau)$ of Lecture 5. The sectors of the chiral algebra *are* a basis for a Hilbert space of a three-dimensional theory. Level 1 gives $\dim\mathcal{V} = 2$: the two characters of Example 5.1, the two-dimensional space on which the S and T matrices of (5.6) were already acting. We wrote down that action a chapter ago without a home for it. Here is the home.

For the mapping class group now acts *on a vector space*, which is what representations require. A diffeomorphism of Σ is a bordism from Σ to itself—the mapping cylinder—so functoriality hands it a linear map on $\mathcal{V}(\Sigma)$, and for the torus the mapping class group is $SL(2,\mathbb{Z})$. The Kac–Peterson matrices are that representation:

$$SL(2,\mathbb{Z}) \;\longrightarrow\; GL(\mathcal{V}(\mathsf{T}^2)), \qquad S \longmapsto S_{jj'}, \quad T \longmapsto T_{jj'}.$$

The theorem of Lecture 5—that the characters span a finite-dimensional $SL(2,\mathbb{Z})$ representation—is now a *corollary of functoriality* in three dimensions, exactly as *partition function = trace* was a corollary of functoriality in two. And the phases: the mapping cylinder construction is only projectively well defined (the framing anomaly, the metric-dependence Chern–Simons was supposed not to have and does, mildly, at one loop). The $\mathrm{e}^{-\pi \mathrm{i}/12}$ we have been carrying since §5.4 is a framing anomaly. It was never a technicality. It was three-dimensional physics, showing through.

[3] E. Witten, *Quantum field theory and the Jones polynomial*, Comm. Math. Phys. **121** (1989); the rigorous construction is Reshetikhin–Turaev, and Turaev's *Quantum Invariants of Knots and 3-Manifolds* is the standard reference.

For the curious: the circle, and what a modular tensor category *is*

Run the ladder down one more rung. If Z_3 assigns a vector space to a surface, what does it assign to a circle—to the codimension-two object? A *category.* And that category is the modular tensor category: its objects are the integrable modules $L(\Lambda)$ of Lecture 5 (the Wilson-line labels of Chern–Simons: a circle in Bord_3 is a line one can colour, and the colours are the sectors); its tensor product is the fusion of those sectors, computed by the Verlinde formula

$$N_{ij}^{\ k} = \sum_m \frac{S_{im}\, S_{jm}\, S^*_{km}}{S_{0m}},$$

the fusion multiplicities *computed from* the Kac–Peterson S-matrix of (5.5); its braiding is the linking of Wilson lines in three dimensions, which is why the category is braided and why the 2d chiral algebra had a monodromy at all; and its S and T data are the mapping-class-group action just constructed. Verlinde's formula is not an outside fact about CFT to be cited. It is the statement that a certain category built from Chern–Simons has an associative fusion, and it holds because $\mathcal{V}$ is functorial: cut the torus two ways, count states both ways, and the counts must agree.

A theory that assigns categories to circles, vector spaces to surfaces, and numbers to 3-manifolds is called **extended**, and the codomain is not Vect but a 2-*category*—the level of structure at which our five costumes finally sit in a single frame. We name it and stop. The reader who wants the full ladder, in every dimension at once, wants Lurie's classification of extended topological field theories.

6.5 The ladder

Stand back. What the last two sections did was not add a chapter to the story; it added an *axis* to it.

assigned to	is	in these notes
a closed 3-manifold	a number	(Chern–Simons invariants; not our business)
a surface Σ	a vector space $\mathcal{V}(\Sigma)$	the *blocks* of L4–L5
a circle S^1	a category	the *sectors*: the MTC

Our entire series lived on one rung. $Z(\mathsf{T}^2_\tau)$ is a number—and a number is what $\mathrm{Bord}_2 \to \mathrm{Vect}$ produces, which is why §6.2 could prove its theorem and why §6.3 could not extend it. The blocks were always one rung up, which is why they had to be smuggled in as "holomorphic functions" with anomalous phases and an $SL(2,\mathbb{Z})$ action nobody could source. Two dimensions was where the ground was solid, as promised in §6.1; it was not where the whole building stood.

The single sentence of these notes therefore wants one amendment. *Assign spaces to space and maps to spacetime; the partition function is what the functor does to a closed spacetime*—true, and it is the theorem of §6.2. But the partition function is a shadow. The object that casts it is a vector of blocks in a finite-dimensional space carrying a projective action of the mapping class group, and to see that object one must stand one dimension up and read the trace as what it always was: not a number, but the pairing that turns a state space into one.

Remark. The student who began the series knowing only one of the five dialects can now say what *was* being counted, and where. A partition function is a weighted trace (L1–L2), which is a path integral on a closed spacetime (L3), which is a graded character (L4–L5), which is the value of a functor on a closed bordism (§6.2)—and which is, finally, the pairing of a chiral block with its conjugate inside a Chern–Simons Hilbert space one dimension up. The Boltzmann factor of Lecture 1 weights the states; the trace of Lecture 2 sums them; the torus of Lecture 4 closes them; and the category of §6.4 is what was doing the counting all along.

Chapter 7

Synthesis: the Rosetta Stone

The whole series condenses into one dictionary. Read each row across: it is the same concept spoken in six languages.

Concept	**Stat Mech**	**Quantum Stats**	**QFT**	**2d CFT**	**Rep / VOA**	**Categ.**
The object Z	$\sum_i \mathrm{e}^{-\beta E_i}$	$\mathrm{Tr}\ \mathrm{e}^{-\beta H}$	$\int \mathcal{D}\phi\, \mathrm{e}^{-S_E}$	$\mathrm{Tr}\ q^{L_0-\frac{c}{24}}\bar{q}^{\bar{L}_0-\frac{c}{24}}$	$\sum_n \dim V_n\, q^n$	Z(closed mfd)
Inverse temp.	$\beta = 1/k_B T$	β	length of time circle	$\mathrm{Im}\,\tau$	formal q	glued cylinder
Conjugate potentials	$\beta\,(E),\ \mu\,(N)$	fugacity $z = \mathrm{e}^{\beta\mu}$	sources / fields	$\tau\,(L_0),\ y\,(J_0)$	τ, Cartan vars	weights
Energy	E_i	spectrum of H	action S_E	$L_0, \bar{L}_0$	grade n	—
Boltzmann weight	$\mathrm{e}^{-\beta E_i}$	$\mathrm{e}^{-\beta H}$	e^{-S_E}	$q^{L_0-c/24}$	$q^{\deg}$	morphism
Charges $\to$ weights	add. $\Rightarrow$ factor.	trace over $\otimes$	local action	grade adds, q's mult.	character of $(\mathrm{ch}, +)$	functoriality
Vacuum subtr.	$\frac{1}{2}\hbar\omega$	—	$\zeta(-1) = -\frac{1}{12}$	$-c/24$ in η	$L_0 - c/24$ shift	—
Sign grading	—	$(-1)^F$	antiperiodicity	spin structure	R-sector $(-1)^F$	—
"Partition" sense	prob. partition	integer partitions $p(N)$	—	—	weight mult.	—
Self-consistency	—	basis indep.	sewing	**modular inv.**	Kac–Peterson / Zhu	gluing = comp.
Boson example	$\frac{1}{2\sinh}$	$\prod \frac{1}{1-q^n}$	scalar det.	$\frac{1}{\sqrt{\tau_2}\lvert\eta\rvert^2}$	Heisenberg VOA	—
Fermion example	$1 + \mathrm{e}^{-\beta\epsilon}$	$\prod(1+q^n)$	Dirac det.	θ/η, spin str.	Clifford super-VOA	—

The single sentence

> *In every column, Z is a weighted trace; reading the trace as a graded dimension makes it a character; demanding that the trace not depend on how spacetime is cut makes it modular—and that is the whole subject.*

Where to go next

- **Statistical and quantum statistics:** Kardar, *Statistical Physics of Particles*; Pathria & Beale, *Statistical Mechanics*; Landau & Lifshitz, vol. 5. For the maximum-entropy derivation of the Boltzmann factor, Jaynes (1957).

- **Thermal and Euclidean field theory:** Kapusta & Gale, *Finite-Temperature Field Theory*; Le Bellac, *Thermal Field Theory*.
- **2d CFT:** Di Francesco, Mathieu & Sénéchal, *Conformal Field Theory*; Polchinski, *String Theory* vol. 1 (the torus partition function, spin structures).
- **Affine Lie algebras and characters:** Kac, *Infinite-Dimensional Lie Algebras*; Kac & Peterson on modular transformations.
- **Vertex operator algebras:** Frenkel & Ben-Zvi, *Vertex Algebras and Algebraic Curves*; Zhu's modular-invariance theorem.
- **Functorial QFT and modular tensor categories:** Atiyah and Segal on functorial (T)QFT; Bakalov & Kirillov, *Lectures on Tensor Categories and Modular Functors*; Turaev. For the elliptic genus, Witten's work on elliptic genera and the index.