Module 2 — Postulates

Quantum Foundations

hbar.university

The State Postulate

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.1

This is the first injection of physical content into the linear-algebraic substrate of Module 1. Up to this point we have built a complex Hilbert space — a mathematical object. The state postulate identifies a specific subset of that object with the physical states of a quantum system.

Formal Statement

To every isolated physical system there is associated a complex separable Hilbert space $\mathcal{H}$, called the **state space** of the system.

The states of the system are represented by **rays** in $\mathcal{H}$: equivalence classes

$$[|\psi\rangle] = \{\alpha|\psi\rangle : \alpha \in \mathbb{C}, \alpha \neq 0\}$$

with $|\psi\rangle \in \mathcal{H}$ nonzero.

By convention, one fixes a representative of unit norm:

$$\||\psi\rangle\| = \sqrt{\langle\psi|\psi\rangle} = 1.$$

Even after normalization, the **global phase** is not determined: $|\psi\rangle$ and $e^{i\theta}|\psi\rangle$ represent the same physical state for any $\theta \in \mathbb{R}$.

Intuition

A quantum state is not a vector. It is a **direction** in Hilbert space — a one-dimensional subspace, with magnitude and global phase stripped away.

The reason is the Born rule (Node 2.3): all measurable predictions take the form $|\langle\phi|\psi\rangle|^2$, and this quantity is invariant under $|\psi\rangle \mapsto e^{i\theta}|\psi\rangle$. Two vectors that differ only by an overall complex factor produce identical measurement statistics, so they cannot represent physically distinct situations.

What survives the quotient by global phase is the **projective Hilbert space** $\mathbb{P}\mathcal{H}$. This — not $\mathcal{H}$ itself — is the true state space of a quantum system.

Crucial caveat: the equivalence is by *global* phase only. **Relative** phases between components of a superposition are physical and observable (interference depends on them).

Structural Role

This postulate fixes the ontology of the rest of the course:

- **Observables** (Node 2.2) will be Hermitian operators on $\mathcal{H}$.
- **Measurement outcomes** (Node 2.3) will be eigenvalues, with probabilities determined by inner products.
- **Time evolution** (Node 2.4) will be unitary, hence ray-preserving.
- **Composite systems** (Node 2.5) will live in tensor products of state spaces.
- **Mixed states** (Node 5.1) will generalize rays to density operators when statistical ensembles are introduced.

Every other postulate is calibrated against this one.

Mathematical Development

The projective Hilbert space

The projective Hilbert space is the quotient

$$\mathbb{P}\mathcal{H} = (\mathcal{H} \setminus \{0\}) / \sim$$

where $|\psi\rangle \sim |\phi\rangle \iff |\phi\rangle = \alpha|\psi\rangle$ for some $\alpha \in \mathbb{C}^*$.

For $\mathcal{H} = \mathbb{C}^{n+1}$, this is the complex projective space $\mathbb{CP}^n$. For the qubit, $\mathcal{H} = \mathbb{C}^2$ and $\mathbb{CP}^1$ is topologically the 2-sphere — the **Bloch sphere**.

Why rays, not vectors

Suppose two states $|\psi\rangle, e^{i\theta}|\psi\rangle$ were physically distinct. Then there would have to exist some measurement procedure that distinguishes them. But by the Born rule (anticipating 2.3), the probability of any outcome $|\phi\rangle$ is

$$|\langle\phi|e^{i\theta}\psi\rangle|^2 = |e^{i\theta}|^2\,|\langle\phi|\psi\rangle|^2 = |\langle\phi|\psi\rangle|^2.$$

The two vectors are operationally indistinguishable, so the postulate identifies them.

Worked Example

A qubit on the Bloch sphere

A general qubit state can be written

$$|\psi\rangle = \cos(\theta/2)\,|0\rangle + e^{i\varphi}\sin(\theta/2)\,|1\rangle$$

with $\theta \in [0, \pi]$, $\varphi \in [0, 2\pi)$. This parameterization eliminates global phase by convention (the coefficient of $|0\rangle$ is real and non-negative).

The map $|\psi\rangle \mapsto (\theta, \varphi)$ sends qubit states bijectively onto the unit sphere $S^2 \subset \mathbb{R}^3$ — the Bloch sphere.

Key facts:

- Each point on the sphere is one quantum state (one ray).
- Antipodal points are orthogonal vectors in $\mathcal{H}$ (e.g., $|0\rangle$ and $|1\rangle$ are at the north and south poles).
- The interior of the ball corresponds to mixed states (introduced in Node 5.1).

Cross-Links

- **Linear Algebra:** Vector Spaces, Quotient Spaces.
- **Statistics:** Classical state = point in a sample space. Quantum state = ray in a Hilbert space. The structural difference is what generates all subsequent quantum-vs-classical distinctions.
- **VQA:** Parameterized quantum circuits prepare rays $|\psi(\theta)\rangle$; the cost function is a real-valued function on the projective Hilbert space.
- **Quantum Information:** Qubit states, qudit states, computational basis.

Structural Compression

Invariant: Physical states are equivalence classes under global phase, not raw Hilbert vectors.

Mechanism: Quotient of $\mathcal{H} \setminus \{0\}$ by complex scalar multiplication.

Cross-System Echo: Probability distributions in classical statistics live on the simplex (a real convex set); quantum states live on the projective Hilbert space (a complex projective manifold). The non-convexity of pure states is the structural marker that quantum mechanics is irreducibly different.

Exercises

1. Show that the relation $|\psi\rangle \sim \alpha|\psi\rangle$ (for $\alpha \in \mathbb{C}^*$) is an equivalence relation on $\mathcal{H} \setminus \{0\}$.
2. Verify that the Bloch sphere parameterization $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$ covers every qubit state exactly once.
3. Show that two qubit states are orthogonal in $\mathbb{C}^2$ if and only if their Bloch vectors are antipodal.
4. Argue why the relative phase between $|0\rangle$ and $|1\rangle$ in a superposition $\alpha|0\rangle + \beta|1\rangle$ is physical, even though the global phase is not.

The Observable Postulate

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.2

The state postulate (2.1) said *what a state is*. This postulate says *what we are allowed to ask about a state*. Together they fix the ontology of measurement before the Born rule (2.3) supplies the statistics.

Formal Statement

To every measurable physical quantity (an **observable**) of a quantum system is associated a self-adjoint (Hermitian) operator on the system's Hilbert space:

$$A : \mathcal{H} \to \mathcal{H}, \qquad A = A^\dagger.$$

The possible outcomes of a measurement of the observable are the elements of the **spectrum** of A:

- For a bounded operator in finite dimensions, the spectrum is the set of real eigenvalues $\{\lambda_i\}$ of the spectral decomposition $A = \sum_i \lambda_i P_i$ (Node 1.6).
- For an unbounded self-adjoint operator on an infinite-dimensional Hilbert space, the spectrum is a closed subset of $\mathbb{R}$ described by a projection-valued measure E_A, and the observable is realized as $A = \int \lambda \, dE_A(\lambda)$.

In finite dimensions, Hermiticity and self-adjointness are the same condition. In infinite dimensions they differ: self-adjointness requires both symmetry ($\langle \phi, A\psi \rangle = \langle A\phi, \psi \rangle$ on the common domain) and equality of domains $\mathrm{dom}(A) = \mathrm{dom}(A^\dagger)$. Only genuinely self-adjoint operators admit a spectral measure, and only these are legitimate observables.

Intuition

An observable is encoded by an operator because measurement in quantum mechanics is non-commutative: the result of measuring "A then B" differs in general from "B then A", and operators are the mathematical objects whose composition is non-commutative.

Hermiticity is chosen because:

1. **Measurement outcomes are real numbers.** The eigenvalues of a Hermitian operator are real (Node 1.6). Complex eigenvalues would correspond to unobservable quantities.
2. **Expectation values are real.** For any state $|\psi\rangle$, the number $\langle \psi | A | \psi \rangle$ is real if and only if A is Hermitian. Reality of expectation values is a necessary condition for A to represent an empirical mean.
3. **Spectral decomposition exists.** Hermiticity is the minimal algebraic condition under which the spectral theorem guarantees an orthogonal decomposition of Hilbert space into eigenspaces — which is what the Born rule (2.3) and projective measurement (3.1) consume.

Non-Hermitian operators still appear in quantum mechanics — raising and lowering operators, annihilation operators, Kraus operators — but they describe *transformations between states*, not measurable quantities. Amplitudes, not outcomes.

Structural Role

The observable postulate is the hinge that turns the spectral theorem (1.6) into physics.

- It dictates that all operators appearing in the Born rule (2.3) come equipped with real spectra and orthogonal eigenspaces.
- It licenses the correspondence "operator eigenvalues $\leftrightarrow$ possible outcomes."
- It sets up the notion of **compatible observables**: observables whose operators commute can be measured jointly; those whose operators do not cannot (Node 3.5).
- It is the structural origin of uncertainty relations — which are statements about pairs of non-commuting observables.

Downstream, it also sets up the Hamiltonian as the distinguished observable that generates time evolution (Node 2.4).

Mathematical Development

Expectation and variance

For a state $|\psi\rangle$ and an observable A, the **expectation value** is

$$\langle A\rangle_\psi := \langle\psi|A|\psi\rangle.$$

With spectral decomposition $A = \sum_i \lambda_i P_i$,

$$\langle A\rangle_\psi = \sum_i \lambda_i \langle\psi|P_i|\psi\rangle = \sum_i \lambda_i \, \|P_i|\psi\rangle\|^2.$$

The right-hand side is a weighted sum of eigenvalues with non-negative weights that sum to one — a genuine probabilistic expectation, as the Born rule (2.3) will confirm.

The **variance** of A in $|\psi\rangle$ is

$$(\Delta A)^2_\psi := \langle A^2\rangle_\psi - \langle A\rangle^2_\psi = \langle\psi|(A - \langle A\rangle_\psi I)^2|\psi\rangle.$$

A state is an **eigenstate** of A if and only if its variance vanishes: the outcome is deterministic.

Commuting observables and joint measurability

Two self-adjoint operators A, B are **compatible** if they commute: $[A, B] = 0$.

Theorem. Two Hermitian operators on a finite-dimensional Hilbert space commute if and only if they are simultaneously diagonalizable in a common orthonormal basis.

The eigenvectors of such a common diagonalization are labelled by joint eigenvalue pairs (λ, μ), and each such pair is a possible outcome of a simultaneous measurement of A and B. Non-commuting observables have no common eigenbasis and cannot be jointly measured with definite outcomes — the structural origin of quantum indeterminacy.

Robertson uncertainty relation

For self-adjoint A and B and any state $|\psi\rangle$,

$$\Delta A \cdot \Delta B \geq \tfrac{1}{2}|\langle[A, B]\rangle_\psi|.$$

The commutator is the algebraic origin of the uncertainty lower bound. Node 3.5 will revisit this; here we note only that it is a direct consequence of Cauchy–Schwarz applied to the state-dependent inner product.

Functional calculus

Because observables are Hermitian, any real-valued function of an observable is again an observable:

$$f(A) := \sum_i f(\lambda_i) P_i \quad \text{(for real } f\text{)}.$$

This is used constantly: A^2, $e^{-\beta A}$, and so on. The functional calculus is what makes "expectation of the square of A" a well-defined observable.

Worked Example

Pauli operators as qubit observables

The three Pauli operators X, Y, Z (Node 1.5) are Hermitian, have spectrum $\{+1, -1\}$, and constitute the basic observables of a qubit.

- **Measuring Z:** outcomes ± 1 corresponding to eigenvectors $|0\rangle, |1\rangle$ — the computational-basis measurement.
- **Measuring X:** outcomes ± 1 corresponding to eigenvectors $|+\rangle, |-\rangle$ — the Hadamard-basis measurement.
- **Measuring Y:** outcomes ± 1 corresponding to eigenvectors $|\pm i\rangle = \frac{1}{\sqrt{2}}(|0\rangle \pm i|1\rangle)$.

Non-commutativity. Since $[X, Z] = -2iY \neq 0$, X and Z are not jointly measurable. There is no state with simultaneously definite values of both. The Robertson uncertainty relation gives

$$\Delta X \cdot \Delta Z \geq |\langle Y \rangle_\psi|,$$

which is saturated when $|\psi\rangle$ is a Y-eigenstate.

Position and momentum on $L^2(\mathbb{R})$

For a particle on a line, $\mathcal{H} = L^2(\mathbb{R})$. The position and momentum observables are

$$(X\psi)(x) = x\psi(x), \qquad (P\psi)(x) = -i\hbar \frac{d\psi}{dx}.$$

Both are unbounded self-adjoint operators, defined on dense domains. Their spectra are the entire real line $\sigma(X) = \sigma(P) = \mathbb{R}$; neither has normalizable eigenvectors, so "position eigenstates" $\delta(x - x_0)$ and "momentum eigenstates" $e^{ipx/\hbar}$ are formal distributions, not Hilbert-space vectors (Node 1.3).

The canonical commutation relation

$$[X, P] = i\hbar I$$

together with Robertson's inequality gives the Heisenberg uncertainty relation

$$\Delta X \cdot \Delta P \geq \frac{\hbar}{2}.$$

Number operator on Fock space

For a single bosonic mode, the Hilbert space is $\mathcal{H} = \bigoplus_{n=0}^{\infty} \mathbb{C}|n\rangle$ — the Fock space. The **number operator** N is defined by $N|n\rangle = n|n\rangle$, with spectrum $\{0, 1, 2, \ldots\}$. It is Hermitian and unbounded. The Hamiltonian of the harmonic oscillator $H = \hbar\omega(N + \frac{1}{2})$ is a real-linear function of N; its eigenvalues are the quantized energy levels.

Cross-Links

- **Linear Algebra:** Self-adjoint operators and the spectral theorem.
- **Statistics:** Random variables are the classical analogue of observables; the structural difference is that classical random variables commute (they are functions on a common sample space) and quantum observables need not.
- **VQA:** The cost-defining Hamiltonian is an observable; variational estimates are expectation values of this observable in parameterized states.
- **Quantum Information:** Measurement bases correspond to choices of observable; all single-qubit measurements reduce to rotations of Z by appropriate unitaries.
- **Quantum Thermodynamics:** Energy is the observable represented by the Hamiltonian; entropy is a function of the density operator, not itself an observable.

Structural Compression

Invariant: Observables are self-adjoint operators; their spectra are real sets of possible outcomes.

Mechanism: Hermiticity guarantees real eigenvalues, real expectation values, and orthogonal eigenspace decomposition — all three requirements of a physical measurement.

Cross-System Echo: Classical random variables on a common sample space are all jointly measurable because they are functions on the same underlying set. Quantum observables live as operators, which need not commute, and their failure to commute is exactly what makes quantum probability a genuine generalization of classical probability rather than an elaborate special case.

Exercises

1. Prove that the expectation value $\langle\psi|A|\psi\rangle$ is real if and only if $A = A^\dagger$.
2. Show that a state $|\psi\rangle$ has zero variance $\Delta A = 0$ if and only if $|\psi\rangle$ is an eigenstate of A.
3. Prove the Robertson uncertainty relation $\Delta A \cdot \Delta B \geq \frac{1}{2}|\langle[A, B]\rangle_\psi|$. (Hint: apply Cauchy–Schwarz to the vectors $(A - \langle A\rangle)|\psi\rangle$ and $(B - \langle B\rangle)|\psi\rangle$.)
4. Compute $\langle X\rangle, \langle Z\rangle, \Delta X, \Delta Z$ for the qubit state $|\psi\rangle = \cos(\theta/2)|0\rangle + \sin(\theta/2)|1\rangle$. Verify that $\Delta X \cdot \Delta Z \geq |\langle Y\rangle|$.
5. Show that two Hermitian operators on a finite-dimensional Hilbert space commute if and only if they admit a common orthonormal eigenbasis.
6. Verify the canonical commutation relation $[X, P] = i\hbar I$ by direct calculation on a smooth test function $\psi(x) \in L^2(\mathbb{R})$.
7. Give an example of a non-Hermitian operator used in quantum mechanics, and explain what role it plays. Argue why it is not itself an observable.

The Born Rule

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.3

The Born rule is the bridge from the quantum formalism to empirical reality. It is the only place where probabilities enter the postulates — and the only place where the formalism contacts measurement statistics. Remove it and everything in Modules 1 and 2 remains mathematically consistent and physically inert.

Formal Statement

Let A be an observable with spectral decomposition

$$A = \sum_i \lambda_i P_i,$$

and let $|\psi\rangle \in \mathcal{H}$ be a normalized state ($\|\psi\| = 1$). When A is measured:

1. **Outcome probabilities.** The probability of obtaining outcome λ_i is

$$p(\lambda_i) = \langle\psi|P_i|\psi\rangle = \|P_i|\psi\rangle\|^2.$$

2. **Post-measurement state (projection postulate).** Conditional on obtaining outcome λ_i, the state immediately after measurement is

$$|\psi'\rangle = \frac{P_i|\psi\rangle}{\sqrt{p(\lambda_i)}}.$$

The post-measurement state lies entirely in the eigenspace $\mathrm{ran}(P_i)$, so an immediately repeated measurement of A yields the same outcome with certainty.

If λ_i is non-degenerate with eigenvector $|\lambda_i\rangle$, the rule specializes to

$$p(\lambda_i) = |\langle\lambda_i|\psi\rangle|^2, \qquad |\psi'\rangle = |\lambda_i\rangle.$$

The first formula — squared amplitude equals probability — is the colloquial "Born rule." The projection step is sometimes called the "collapse postulate"; its interpretive status is revisited in Module 6.

Intuition

A state $|\psi\rangle$ is a superposition of eigenstates of whichever observable you choose to measure. The Born rule reads off, from the amplitude of each eigenstate component, the probability that the measurement outcome will be the corresponding eigenvalue. "The state is α parts eigenstate-1 and β parts eigenstate-2" translates to "outcome 1 occurs with probability $|\alpha|^2$, outcome 2 with probability $|\beta|^2$."

The squared modulus is essential. It is what discards global phase (making physical states rays, per Node 2.1), preserves total probability (via the resolution of identity), and aligns complex amplitudes with real probabilities.

The projection postulate is the statement that, having observed a particular outcome, the state is now compatible with that outcome — it has been projected into the corresponding eigenspace. This is the single step in quantum mechanics that is not unitary, and it is the source of the measurement problem (Module 6).

Structural Role

The Born rule is the empirical interface of the theory. Every laboratory number that has ever been measured in a quantum experiment — every photon count, every neutron polarization, every qubit readout — is a sample from a Born-rule probability distribution.

- It consumes the spectral decomposition of observables (Node 1.6, 2.2).
- It is the foundation of projective measurement (Node 3.1), POVMs (Node 3.2), and weak measurement (Node 3.5).
- It is the ingredient that extends to density matrices (Node 5.1) via $p(\lambda_i) = \mathrm{Tr}(P_i \rho)$.
- It is what variational quantum algorithms (VQA course) sample from: every estimate of $\langle\psi(\theta)|H|\psi(\theta)\rangle$ is a finite-shot average over Born-rule samples.

Removing the Born rule leaves the rest of the postulates mutually consistent but empirically silent.

Mathematical Development

Total probability sums to one

From the resolution of identity $\sum_i P_i = I$ (spectral theorem, Node 1.6):

$$\sum_i p(\lambda_i) = \sum_i \langle\psi|P_i|\psi\rangle = \langle\psi| \left(\sum_i P_i\right) |\psi\rangle = \langle\psi|I|\psi\rangle = \|\psi\|^2 = 1.$$

This is why normalization matters and why the spectral theorem is the structural prerequisite. Total probability is conserved as an algebraic identity.

Expectation value

The expectation value of an observable equals the Born-weighted mean of its eigenvalues:

$$\langle A\rangle_\psi = \sum_i \lambda_i p(\lambda_i) = \sum_i \lambda_i \langle\psi|P_i|\psi\rangle = \langle\psi|A|\psi\rangle.$$

The first and last expressions are equal by the spectral decomposition. This identification — "expectation value equals $\langle\psi|A|\psi\rangle$" — is a *theorem*, not an additional postulate, once the Born rule and the spectral theorem are in place.

Variance and uncertainty

The variance of A in the state $|\psi\rangle$ is

$$(\Delta A)^2 = \sum_i (\lambda_i - \langle A\rangle)^2 p(\lambda_i) = \langle A^2\rangle - \langle A\rangle^2 = \langle\psi|(A - \langle A\rangle I)^2|\psi\rangle.$$

Zero variance $\iff$ $|\psi\rangle$ is an eigenstate of A $\iff$ the outcome is deterministic.

Degenerate eigenvalues

When λ_i has an eigenspace of dimension > 1, P_i is a projection onto that subspace. The outcome probability $\langle\psi|P_i|\psi\rangle$ depends only on the projection of $|\psi\rangle$ onto the eigenspace, not on the internal direction. The post-measurement state is $P_i|\psi\rangle/\sqrt{p(\lambda_i)}$, which is generally not an eigenvector of A alone — it is merely an element of the correct eigenspace.

Sequential measurements

Consider measuring A, obtaining λ_i, and then measuring $B = \sum_j \mu_j Q_j$. The joint probability is

$$p(\lambda_i, \mu_j) = \langle\psi|P_i Q_j P_i|\psi\rangle,$$

which in general is *not* the classical joint probability $p(\lambda_i)p(\mu_j|\lambda_i)$ with commuting conditionals. If $[A, B] = 0$, the projections commute and sequential measurements behave classically. If not, the order of measurement matters — the structural origin of quantum contextuality (Node 10.2).

Worked Example

Pauli Z on $|+\rangle$

Prepare the qubit state

$$|+\rangle = \tfrac{1}{\sqrt{2}}(|0\rangle + |1\rangle),$$

and measure the observable Z, with spectral decomposition $Z = (+1)|0\rangle\langle 0| + (-1)|1\rangle\langle 1|$.

Outcome probabilities:

$$p(+1) = |\langle 0|+\rangle|^2 = \left|\tfrac{1}{\sqrt{2}}\right|^2 = \tfrac{1}{2}, \qquad p(-1) = |\langle 1|+\rangle|^2 = \tfrac{1}{2}.$$

Post-measurement states: if the outcome is $+1$, the state collapses to $|0\rangle$; if -1, to $|1\rangle$.

Expectation value:

$$\langle Z\rangle_+ = (+1)\cdot\tfrac{1}{2} + (-1)\cdot\tfrac{1}{2} = 0 = \langle +|Z|+\rangle.$$

Variance:

$$(\Delta Z)^2_+ = \langle Z^2\rangle_+ - \langle Z\rangle^2_+ = 1 - 0 = 1.$$

The state $|+\rangle$ is maximally uncertain in Z, minimally uncertain (variance zero) in X — the tradeoff enforced by non-commutativity.

Repeated measurement

Prepare $|+\rangle$, measure Z, suppose the outcome is $+1$. The state is now $|0\rangle$. Measure Z again: the outcome is $+1$ with probability $|\langle 0|0\rangle|^2 = 1$. Deterministic repetition — the structural content of the projection postulate.

If instead, after measuring Z and obtaining $+1$, we measure X, we get ± 1 with probability $|\langle \pm|0\rangle|^2 = \frac{1}{2}$. The Z measurement has erased the state's X information.

Cross-Links

- **Linear Algebra:** The spectral theorem supplies $\sum_i P_i = I$, which is the algebraic source of "probabilities sum to one."
- **Statistics:** Probability measures on outcome spaces are the classical analogue. The Born rule is a non-commutative generalization — the probability space is state-dependent and observable-dependent.
- **VQA:** Every cost-function estimate is a finite-shot sample from a Born distribution, introducing shot-noise as a fundamental error source.
- **Quantum Information:** Distinguishability of quantum states is bounded by Born-rule overlap $|\langle\phi|\psi\rangle|^2$; non-orthogonal states are not perfectly distinguishable.
- **Module 5 (Decoherence):** The density-matrix form $p(\lambda_i) = \mathrm{Tr}(P_i\rho)$ generalizes the pure-state Born rule to statistical mixtures.
- **Module 6 (Interpretations):** The projection postulate is the single non-unitary step in the postulates; its interpretive status is the measurement problem.

Structural Compression

Invariant: Squared amplitudes are probabilities, and spectral projections resolve the identity.

Mechanism: $p(\lambda_i) = \langle\psi|P_i|\psi\rangle$, with total probability enforced by $\sum_i P_i = I$.

Cross-System Echo: Classical expectation is $\int X\,dP$ with dP a probability measure on a sample space. Quantum expectation is $\langle\psi|A|\psi\rangle$ with the "measure" depending on both the state and the observable. The structural role is identical; the non-commutativity is what is new.

Exercises

1. Verify that the Born rule gives total probability one by direct computation: $\sum_i \langle\psi|P_i|\psi\rangle = 1$ for any unit vector $|\psi\rangle$.
2. Derive the identity $\langle A\rangle_\psi = \sum_i \lambda_i p(\lambda_i)$ from the Born rule and the spectral decomposition.
3. For $|\psi\rangle = \cos(\theta/2)|0\rangle + e^{i\varphi}\sin(\theta/2)|1\rangle$, compute the Born probabilities for a Z-measurement and for an X-measurement. Explain why the answer for Z depends only on θ while the answer for X depends on both θ and φ.
4. Prepare $|0\rangle$ and measure X. What are the probabilities of each outcome and the post-measurement state in each case?
5. Show that if $[A, B] = 0$, then sequential measurements of A and B produce the same statistics as a single joint measurement (in a common eigenbasis). Show by example that this fails when $[A, B] \neq 0$.
6. Derive the generalization $p(\lambda_i) = \mathrm{Tr}(P_i\rho)$ of the Born rule from the pure-state version, assuming $\rho = \sum_k p_k |\psi_k\rangle\langle\psi_k|$ is a statistical mixture of pure states. (This previews Node 5.1.)
7. The Born rule's squared modulus is sometimes called "the mysterious quadratic." Write a short paragraph arguing, structurally, why linear amplitudes ($p \propto |\alpha|$) and cubic amplitudes ($p \propto |\alpha|^3$) are both incompatible with the inner-product structure of Hilbert space.

Unitary Time Evolution

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.4

This postulate gives the formalism its dynamics. The state postulate (2.1) and the observable postulate (2.2) define a static theory of states and questions; the Born rule (2.3) turns it into a theory of measurement statistics. Time evolution is what makes it a theory of how states change between measurements.

Formal Statement

The time evolution of an isolated quantum system is generated by a self-adjoint operator H — the **Hamiltonian**, identified with the system's total energy observable — via the **Schrödinger equation**:

$$i\hbar \frac{d}{dt}|\psi(t)\rangle = H|\psi(t)\rangle.$$

Equivalently, evolution is implemented by a two-parameter family of **unitary operators** $U(t_2, t_1) : \mathcal{H} \to \mathcal{H}$:

$$|\psi(t_2)\rangle = U(t_2, t_1)|\psi(t_1)\rangle,$$

satisfying $U^\dagger(t_2, t_1)U(t_2, t_1) = I$, the composition law $U(t_3, t_1) = U(t_3, t_2)U(t_2, t_1)$, and the boundary condition $U(t, t) = I$.

For a time-independent Hamiltonian the solution is explicit:

$$U(t) = e^{-iHt/\hbar}, \qquad |\psi(t)\rangle = e^{-iHt/\hbar}|\psi(0)\rangle.$$

For a time-dependent Hamiltonian $H(t)$, the formal solution is the **time-ordered exponential**

$$U(t, 0) = \mathcal{T} \exp\left(-\frac{i}{\hbar}\int_0^t H(s)\, ds\right).$$

Intuition

Unitary evolution is the formal expression of two physical demands:

1. **Probabilities are conserved.** The total probability of all possible measurement outcomes must remain one at every instant; otherwise the theory predicts the disappearance (or creation) of probability mass, which is empirically absurd.
2. **Evolution is linear and deterministic.** The superposition principle survives dynamics: if $|\psi_1(t)\rangle$ and $|\psi_2(t)\rangle$ are solutions, so is $\alpha|\psi_1(t)\rangle + \beta|\psi_2(t)\rangle$, with the same coefficients at every time.

Unitarity is the unique class of operators on a Hilbert space that is both linear and norm-preserving. It is therefore the only kind of map that can represent the closed-system dynamics demanded by the postulates.

The Hamiltonian plays a double role: as the energy observable (a Hermitian operator whose eigenvalues are the possible energy measurement outcomes) and as the generator of time evolution (whose exponential supplies the unitary). This double role is not a coincidence: it is the quantum version of Noether's theorem for time translation (Node 2.6).

Structural Role

Unitary evolution is the postulate that separates **closed systems** from **open systems**:

- A closed system evolves unitarily under a fixed Hamiltonian, and its state remains pure (Node 5.1).
- An open system coupled to an environment does not evolve unitarily on its own; its effective dynamics are described by completely positive trace-preserving maps (Node 7.3) and generated by Lindblad operators (Node 8.2).

Every structural feature of Module 8 (open systems) is derived by starting from a unitary evolution on a larger system and tracing out part of it. Unitarity is therefore not merely the dynamics of closed systems; it is the axiom from which all other quantum dynamics are built.

Unitary evolution is also the fundamental resource of the VQA course: quantum circuits are products of unitary gates, each implementing $U(\theta) = e^{-iG\theta}$ for some generator G.

Mathematical Development

Why unitary? The Wigner argument

Suppose evolution $|\psi\rangle \mapsto T|\psi\rangle$ preserves all transition probabilities: $|\langle T\phi|T\psi\rangle|^2 = |\langle\phi|\psi\rangle|^2$ for all $|\phi\rangle, |\psi\rangle$. Wigner's theorem asserts that T must be either unitary (linear, inner-product preserving) or anti-unitary (antilinear, conjugate-inner-product preserving). Continuous time evolution rules out the anti-unitary case by a continuity-from-the-identity argument, leaving unitarity as the unique option.

Stone's theorem

Theorem (Stone). A one-parameter group $U(t)$ of unitary operators on a Hilbert space is strongly continuous if and only if there exists a (possibly unbounded) self-adjoint operator H such that

$$U(t) = e^{-iHt/\hbar}.$$

The operator H is called the **infinitesimal generator** of the group. Stone's theorem is the rigorous statement that "continuous unitary evolution" and "Hermitian Hamiltonian" are two sides of the same structural object.

Probability conservation

If $|\psi(t)\rangle = U(t)|\psi(0)\rangle$ with $U(t)$ unitary, then

$$\langle\psi(t)|\psi(t)\rangle = \langle\psi(0)|U^\dagger U|\psi(0)\rangle = \langle\psi(0)|\psi(0)\rangle,$$

so norms (hence Born-rule total probabilities) are conserved exactly.

The three pictures

Time dependence can be assigned to states or operators:

- **Schrödinger picture.** States evolve: $|\psi_S(t)\rangle = U(t)|\psi(0)\rangle$. Operators are time-independent. This is the default.
- **Heisenberg picture.** States are fixed; operators evolve: $A_H(t) = U^\dagger(t)AU(t)$. The Heisenberg equation reads $\frac{dA_H}{dt} = \frac{i}{\hbar}[H, A_H]$.
- **Interaction picture.** For $H = H_0 + V(t)$, the fast part H_0 is absorbed into the operators and the slow part V drives the remaining evolution of the states. Used for perturbation theory.

Expectation values are identical across pictures: $\langle\psi_S(t)|A|\psi_S(t)\rangle = \langle\psi(0)|A_H(t)|\psi(0)\rangle$.

Ehrenfest's theorem

For any observable A,

$$\frac{d}{dt}\langle A\rangle_{\psi(t)} = \frac{i}{\hbar}\langle [H, A]\rangle_{\psi(t)} + \left\langle \frac{\partial A}{\partial t} \right\rangle.$$

The classical limit recovers Hamilton's equations: $\dot{q} = \partial H/\partial p$, $\dot{p} = -\partial H/\partial q$, with Poisson brackets replacing commutators by $\{\cdot,\cdot\} \leftrightarrow \frac{1}{i\hbar}[\cdot,\cdot]$.

Worked Example

Larmor precession of a qubit

A spin-$\frac{1}{2}$ particle in a static magnetic field along $\hat{z}$ has Hamiltonian

$$H = \frac{\hbar\omega}{2} Z, \qquad \omega = \gamma B,$$

where γ is the gyromagnetic ratio and B the field strength. The time evolution operator is

$$U(t) = e^{-iHt/\hbar} = e^{-i\omega t Z/2} = \cos(\omega t/2) I - i \sin(\omega t/2) Z.$$

Evolution of $|+\rangle$

Initial state: $|\psi(0)\rangle = |+\rangle = \frac{1}{\sqrt{2}}(|0\rangle + |1\rangle)$.

$$|\psi(t)\rangle = U(t)|+\rangle = \tfrac{1}{\sqrt{2}}\left(e^{-i\omega t/2}|0\rangle + e^{+i\omega t/2}|1\rangle\right).$$

Up to an overall phase, this is

$$|\psi(t)\rangle \simeq \tfrac{1}{\sqrt{2}}\left(|0\rangle + e^{i\omega t}|1\rangle\right) = \cos\!\left(\tfrac{\omega t}{2}\right)|+\rangle - i\sin\!\left(\tfrac{\omega t}{2}\right)|-\rangle.$$

Expectation values

Track $\langle X\rangle, \langle Y\rangle, \langle Z\rangle$:

$$\langle X(t)\rangle = \cos(\omega t), \qquad \langle Y(t)\rangle = -\sin(\omega t), \qquad \langle Z(t)\rangle = 0.$$

The Bloch vector $\vec{r}(t) = (\cos\omega t, -\sin\omega t, 0)$ precesses clockwise around the z-axis with angular frequency ω. This is Larmor precession — one of the most directly testable consequences of the Schrödinger equation and the experimental basis of NMR and MRI.

Energy measurement

If instead we measure the energy $H = (\hbar\omega/2)Z$ on $|\psi(t)\rangle$, the outcomes are $\pm\hbar\omega/2$ with probabilities $1/2, 1/2$ — constant in time, because $[H, H] = 0$ makes energy a conserved quantity. Expectation value $\langle H\rangle = 0$ throughout the precession.

Cross-Links

- **Linear Algebra:** Unitary matrices, matrix exponentials, and the operator-valued functional calculus.
- **VQA:** Parameterized circuits are products of unitary gates $U_k(\theta_k) = e^{-iG_k\theta_k}$; gradients are computed via the parameter-shift rule that exploits the unitary structure.
- **Quantum Information:** Quantum gates are unitaries; the universal gate set is a generating set for $U(2^n)$.
- **Quantum Thermodynamics:** Unitary evolution is reversible and conserves von Neumann entropy; the second law arises only when part of the system is coarse-grained away (Module 8).
- **Systems Thinking:** Unitary evolution is the canonical example of a deterministic, reversible, information-preserving dynamical flow.
- **Module 10 (Foundations):** The universal unitarity of closed-system quantum mechanics is in tension with the non-unitary projection postulate — the formal statement of the measurement problem.

Structural Compression

Invariant: Probability conservation, realized as inner-product preservation.

Mechanism: $U^\dagger U = I$, generated via Stone's theorem by a self-adjoint Hamiltonian.

Cross-System Echo: Symplectic flow in classical Hamiltonian mechanics preserves phase-space volume (Liouville's theorem). Unitary flow in quantum mechanics preserves the inner product. Both are the "time-translation symmetry" of their respective state spaces, and both identify the Hamiltonian as the generator.

Exercises

1. Verify that $U(t) = e^{-iHt/\hbar}$ is unitary for any Hermitian H. (Use $(e^{-iHt/\hbar})^\dagger = e^{+iHt/\hbar}$ and commutativity of H with itself.)
2. Derive the Heisenberg equation of motion $dA_H/dt = (i/\hbar)[H, A_H]$ from the definition $A_H(t) = U^\dagger(t)AU(t)$.
3. Prove Ehrenfest's theorem in the Schrödinger picture by differentiating $\langle\psi(t)|A|\psi(t)\rangle$ directly.
4. For $H = (\hbar\omega/2)X$, compute $U(t)$ and apply it to the initial state $|0\rangle$. Sketch the trajectory of the Bloch vector on the sphere.
5. For a two-level system with Hamiltonian $H = \omega_1 X + \omega_2 Z$, find the eigenvalues and eigenvectors of H. Write $U(t)$ in closed form.
6. Show that energy expectation values are constant in time under unitary evolution with a time-independent Hamiltonian. What is the corresponding statement when $H = H(t)$?
7. Explain, structurally, why the projection postulate (Node 2.3) cannot be unitary, and why this discontinuity is the precise location of the measurement problem.

Composite Systems and Tensor Products

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.5

This postulate is the structural origin of entanglement. The tensor product is *not* the Cartesian product, and that single distinction is responsible for everything quantum mechanics has that classical mechanics does not — the exponential state-space growth, Bell nonlocality, quantum computational speedups, and decoherence. Module 4 is devoted to exploring the consequences; here we lay down the structure.

Formal Statement

If a quantum system AB consists of two subsystems A and B with state spaces $\mathcal{H}_A$ and $\mathcal{H}_B$, then the state space of the composite system is the **tensor product**

$$\mathcal{H}_{AB} = \mathcal{H}_A \otimes \mathcal{H}_B.$$

Given orthonormal bases $\{|i\rangle_A\}_{i=1}^{d_A}$ for $\mathcal{H}_A$ and $\{|j\rangle_B\}_{j=1}^{d_B}$ for $\mathcal{H}_B$, an orthonormal basis for $\mathcal{H}_{AB}$ is

$$\{|i\rangle_A \otimes |j\rangle_B : 1 \leq i \leq d_A,\ 1 \leq j \leq d_B\},$$

often abbreviated $|i, j\rangle$ or $|ij\rangle$. Hence

$$\dim \mathcal{H}_{AB} = d_A \cdot d_B.$$

If A is prepared in state $|\psi\rangle_A$ and B in state $|\phi\rangle_B$ independently, the composite state is the **product state**

$$|\psi\rangle_A \otimes |\phi\rangle_B \in \mathcal{H}_{AB}.$$

Product states span $\mathcal{H}_{AB}$ but do not exhaust it: a generic element of $\mathcal{H}_{AB}$ is a linear combination of product states. Such non-product states are called **entangled**, and are the subject of Module 4.

Operators on subsystems extend to operators on the composite via the tensor product of operators:

$$(A \otimes B)(|\psi\rangle_A \otimes |\phi\rangle_B) := (A|\psi\rangle_A) \otimes (B|\phi\rangle_B),$$

extended by linearity to all of $\mathcal{H}_{AB}$. An operator acting on A alone is $A \otimes I_B$; one acting on B alone is $I_A \otimes B$.

Intuition

The structural point is the comparison with classical composite systems.

Structure	Classical (product of sets)	Quantum (tensor of spaces)
Compose state spaces	$S_A \times S_B$	$\mathcal{H}_A \otimes \mathcal{H}_B$
Dimension rule	$\dim(S_A \times S_B) = \dim S_A + \dim S_B$	$\dim(\mathcal{H}_A \otimes \mathcal{H}_B) = \dim \mathcal{H}_A \cdot \dim \mathcal{H}_B$
Generic joint state	factorizable: (a, b)	generally not factorizable
Resource	simple storage	exponential superposition

Two classical bits have $2 + 2 = 4$ degrees of freedom — but they describe one of $2 \times 2 = 4$ possible states. Two quantum bits have $2 \cdot 2 = 4$ **complex amplitudes**, one for each computational basis state, all of which can be simultaneously nonzero. The *amplitude space* of n qubits has dimension 2^n, which is the structural origin of every "quantum advantage" story.

Entanglement is the statement that most vectors in $\mathcal{H}_A \otimes \mathcal{H}_B$ are not products of anything. The tensor product enlarges the state space beyond what "independent preparation" can reach, and the extra room is what Module 4 will study.

Structural Role

Every multi-system construction in the course is built on this postulate:

- **Multi-qubit registers** (Node 7.1) live in $(\mathbb{C}^2)^{\otimes n}$.
- **Bell states and CHSH** (Node 4.4) are the simplest entangled states and violate a classical factorization bound.
- **Partial trace** (Node 5.2) reduces a joint state to a subsystem state and is the structural origin of decoherence and mixed states.
- **Environments** (Node 5.3) are modelled as tensor factors, and Module 8's open-systems dynamics arise from tracing over them.
- **Quantum channels** (Node 7.3) are constructed via Stinespring dilation: every channel is a unitary on a larger tensor-product Hilbert space followed by a partial trace.

Remove the tensor product, or replace it with a direct sum, and none of these exist. The direct sum $\mathcal{H}_A \oplus \mathcal{H}_B$ gives a space of dimension $d_A + d_B$ — which is what you would need if "composite" meant "exclusive alternative" rather than "simultaneous subsystems."

Mathematical Development

Bilinearity

The tensor product $\otimes : \mathcal{H}_A \times \mathcal{H}_B \to \mathcal{H}_A \otimes \mathcal{H}_B$ is bilinear:

$$(\alpha|\psi_1\rangle + \beta|\psi_2\rangle) \otimes |\phi\rangle = \alpha|\psi_1\rangle \otimes |\phi\rangle + \beta|\psi_2\rangle \otimes |\phi\rangle,$$

and similarly in the second argument.

Inner product on the tensor space

The inner product on $\mathcal{H}_A \otimes \mathcal{H}_B$ is determined by

$$\langle \psi_1 \otimes \phi_1 | \psi_2 \otimes \phi_2 \rangle_{AB} := \langle \psi_1 | \psi_2 \rangle_A \, \langle \phi_1 | \phi_2 \rangle_B,$$

extended by sesquilinearity. The product basis $\{|i\rangle \otimes |j\rangle\}$ is then orthonormal, confirming that it is a basis of a Hilbert space of dimension $d_A d_B$.

Operator tensor product

For operators A on $\mathcal{H}_A$ and B on $\mathcal{H}_B$:

$$(A \otimes B)(|\psi\rangle \otimes |\phi\rangle) = (A|\psi\rangle) \otimes (B|\phi\rangle).$$

Key identities:

$$(A_1 \otimes B_1)(A_2 \otimes B_2) = (A_1 A_2) \otimes (B_1 B_2),$$

$$(A \otimes B)^\dagger = A^\dagger \otimes B^\dagger,$$

$$\mathrm{Tr}(A \otimes B) = \mathrm{Tr}(A)\,\mathrm{Tr}(B).$$

Local operations

An operator of the form $A \otimes I_B$ is called **local on** A: it acts on $\mathcal{H}_A$ without touching $\mathcal{H}_B$. Two local operators $A_A \otimes I_B$ and $I_A \otimes B_B$ always commute, regardless of whether A and B would commute on their own subsystem. This is the formal reason why a measurement on one half of an entangled pair does not mechanically disturb the other — the operators commute, so their order does not affect joint statistics. The subtlety of Bell-type experiments (Node 10.1) is that correlations in the outcomes nevertheless exceed what any local hidden-variable model can produce.

Entanglement vs separability — the formal test

A pure state $|\Psi\rangle_{AB}$ is **separable** (equivalently, a product state) if it can be written as $|\psi\rangle_A \otimes |\phi\rangle_B$ for some $|\psi\rangle_A, |\phi\rangle_B$. Otherwise it is **entangled**.

Criterion. Write $|\Psi\rangle_{AB} = \sum_{ij} c_{ij} |i\rangle|j\rangle$ in a product basis, and form the $d_A \times d_B$ matrix C of coefficients. The state is separable if and only if $\mathrm{rank}(C) = 1$. Otherwise it is entangled, and the Schmidt decomposition (Node 4.3) gives its canonical form.

Worked Example

Two qubits

$\mathcal{H}_{AB} = \mathbb{C}^2 \otimes \mathbb{C}^2 \cong \mathbb{C}^4$. Basis:

$$|00\rangle = \begin{pmatrix}1\\0\\0\\0\end{pmatrix}, \quad |01\rangle = \begin{pmatrix}0\\1\\0\\0\end{pmatrix}, \quad |10\rangle = \begin{pmatrix}0\\0\\1\\0\end{pmatrix}, \quad |11\rangle = \begin{pmatrix}0\\0\\0\\1\end{pmatrix}.$$

A product state

Prepare A in $|+\rangle = (|0\rangle + |1\rangle)/\sqrt{2}$ and B in $|0\rangle$. The composite state is

$$|+\rangle_A \otimes |0\rangle_B = \tfrac{1}{\sqrt{2}}(|00\rangle + |10\rangle).$$

Coefficient matrix:

$$C = \tfrac{1}{\sqrt{2}} \begin{pmatrix}1 & 0\\1 & 0\end{pmatrix}, \qquad \mathrm{rank}(C) = 1.$$

Separable — as expected, since it was prepared by independent operations on the two subsystems.

A Bell state (entangled)

The Bell state

$$|\Phi^+\rangle = \tfrac{1}{\sqrt{2}}(|00\rangle + |11\rangle)$$

has coefficient matrix

$$C = \tfrac{1}{\sqrt{2}}\begin{pmatrix}1 & 0\\ 0 & 1\end{pmatrix}, \qquad \mathrm{rank}(C) = 2.$$

No factorization $|\Phi^+\rangle = |\psi\rangle_A \otimes |\phi\rangle_B$ exists. Attempting one leads to the contradiction $\psi_0\phi_0 = 1/\sqrt{2}$, $\psi_0\phi_1 = 0$, $\psi_1\phi_0 = 0$, $\psi_1\phi_1 = 1/\sqrt{2}$: the middle two force $\psi_0 = 0$ or $\phi_1 = 0$ (and similarly on the other pair), which is incompatible with the outer two being nonzero.

Local operator action

The operator $Z \otimes I$ acts on the Bell state as

$$(Z \otimes I)|\Phi^+\rangle = \tfrac{1}{\sqrt{2}}(Z|0\rangle \otimes |0\rangle + Z|1\rangle \otimes |1\rangle) = \tfrac{1}{\sqrt{2}}(|00\rangle - |11\rangle) = |\Phi^-\rangle.$$

A local operation on A transforms the entangled state into another (still entangled) state. The structural fact here is that Bell states form an orthonormal basis of $\mathcal{H}_{AB}$, and local Pauli operators act transitively on this basis.

Cross-Links

- **Linear Algebra:** Tensor products of vector spaces, Kronecker products of matrices, rank of a matrix.
- **Quantum Information:** n-qubit register $(\mathbb{C}^2)^{\otimes n}$ and universal gate sets acting on it.
- **VQA:** Parameterized circuits act on tensor-product spaces; circuit depth and entangling structure determine expressivity.
- **Module 4 (Entanglement):** Full development of entanglement structure, measures, and monogamy.
- **Module 5 (Decoherence):** Partial trace over the environment tensor factor is the operation that turns entanglement with the environment into apparent classical mixtures.
- **Statistics:** Joint distributions on a product sample space correspond to classical (separable) composite states; entanglement has no classical analogue.

Structural Compression

Invariant: $\dim(\mathcal{H}_A \otimes \mathcal{H}_B) = \dim \mathcal{H}_A \cdot \dim \mathcal{H}_B$.

Mechanism: Bilinear extension over a product basis $\{|i\rangle \otimes |j\rangle\}$, with inner product and operators inherited multiplicatively.

Cross-System Echo: Cartesian product of classical state spaces gives additive dimension; tensor product of quantum state spaces gives multiplicative dimension. The exponential gap — 2^n complex amplitudes for n qubits versus $2n$ real coordinates for n classical bits — is exactly the resource that quantum computation, entanglement-based cryptography, and Bell nonlocality each exploit in different ways.

Exercises

1. Verify that the product basis $\{|i\rangle \otimes |j\rangle\}$ is orthonormal with respect to the inner product $\langle \psi_1 \otimes \phi_1 | \psi_2 \otimes \phi_2 \rangle := \langle \psi_1 | \psi_2 \rangle \langle \phi_1 | \phi_2 \rangle$.
2. Let A be a 2×2 matrix and B be a 3×3 matrix. Write out $A \otimes B$ as a 6×6 Kronecker matrix for general entries.
3. Show that $(A \otimes B)(C \otimes D) = (AC) \otimes (BD)$ as operators on $\mathcal{H}_A \otimes \mathcal{H}_B$.
4. Prove that all four Bell states,

$$|\Phi^{\pm}\rangle = \tfrac{1}{\sqrt{2}}(|00\rangle \pm |11\rangle), \qquad |\Psi^{\pm}\rangle = \tfrac{1}{\sqrt{2}}(|01\rangle \pm |10\rangle),$$

are entangled, by computing the rank of each coefficient matrix.

5. Given $|\psi\rangle_{AB} = \frac{1}{2}(|00\rangle + |01\rangle + |10\rangle + |11\rangle)$, determine whether the state is separable or entangled. If separable, write it explicitly as $|\psi\rangle_A \otimes |\phi\rangle_B$.
6. Show that $\mathrm{Tr}(A \otimes B) = \mathrm{Tr}(A)\,\mathrm{Tr}(B)$ using a product basis.
7. Argue, structurally, why n qubits have exponentially more states than n classical bits despite being "the same number of systems," and relate this to the multiplicative dimension rule.

Symmetries and Conservation Laws

Position in Knowledge Graph

System: Quantum Foundations **Structural Domain:** Postulates of Quantum Mechanics **Node:** 2.6

Symmetries are the source of conservation laws in classical and quantum mechanics alike. In the quantum case the structural statement is sharper: a symmetry is a unitary (or anti-unitary) operator that commutes with the Hamiltonian, and its generator is an exactly conserved observable. This node closes Module 2 by connecting the dynamical postulate (2.4) to the algebraic structure of observables (2.2), setting up the appearance of symmetries in thermodynamic equilibria (Module 9) and foundational no-go theorems (Module 10).

Formal Statement

A **symmetry** of a quantum system is a unitary operator $U : \mathcal{H} \to \mathcal{H}$ that commutes with the Hamiltonian:

$$[U, H] = UH - HU = 0.$$

(For time-reversal, anti-unitary operators are also permitted; Wigner's theorem guarantees that every symmetry is either unitary or anti-unitary.)

A **continuous symmetry** is a one-parameter unitary group $U(s) = e^{-iGs}$ with G self-adjoint. The operator G is the **generator** of the symmetry. The commutation condition $[U(s), H] = 0$ for all s is equivalent to

$$[G, H] = 0.$$

Quantum Noether theorem. If $[G, H] = 0$, then the expectation value of G is conserved under time evolution:

$$\frac{d}{dt}\langle G\rangle_{\psi(t)} = 0.$$

More strongly, the full probability distribution of G (not just its mean) is conserved, because G and H are simultaneously diagonalizable and time evolution acts on G-eigenspaces by a phase.

Intuition

If the Hamiltonian does not "see" a symmetry operator, then the symmetry operator does not change under time evolution. The Hamiltonian is the generator of time translation (Node 2.4); the symmetry generator G is the generator of some other transformation (a spatial rotation, a phase shift, a particle exchange). When the two generators commute, the two transformations are independent, and neither flow disturbs the other.

The classical analogue is Noether's theorem: a continuous symmetry of the Lagrangian yields a conserved current. The quantum statement is not just an analogue but a sharpening: in quantum mechanics the conservation is exact (no coarse-graining), and it applies to the full operator spectrum, not just the expectation value.

Symmetries also carve Hilbert space into **irreducible representations** of the symmetry group. A state in a particular irreducible representation carries definite quantum numbers (angular momentum, parity, charge, colour), and these labels are protected by the symmetry: no Hamiltonian respecting the symmetry can mix different irreducible sectors. This is the structural origin of the periodic table, selection rules in spectroscopy, and particle classification in high-energy physics.

Structural Role

Symmetries are the organizing principle of Hilbert space.

- **Invariant subspaces.** A symmetry group decomposes $\mathcal{H}$ into orthogonal invariant subspaces, each carrying an irreducible representation.
- **Good quantum numbers.** Generators that commute with H are simultaneously diagonalizable with H; their eigenvalues label energy levels.
- **Degeneracy.** Multi-dimensional irreducible representations imply multi-dimensional eigenspaces of H — all degeneracies in quantum mechanics trace either to a symmetry or to an accident.
- **Selection rules.** Transition matrix elements between states in different irreducible representations vanish for symmetry-respecting perturbations. This is the origin of "forbidden transitions."
- **Equilibrium ensembles.** In Module 9, Gibbs states take the form $e^{-\beta H}/Z$, which is a function of H alone and therefore commutes with every symmetry of H; this is why conserved charges appear as chemical potentials.

Symmetry reasoning is a labor-saving device in computation and an ontological device in interpretation: it tells you which degrees of freedom are "really there" and which are gauge.

Mathematical Development

Minimum representation theory

A **group representation** is a homomorphism $\rho : G \to U(\mathcal{H})$: each group element g is mapped to a unitary $\rho(g)$ with $\rho(g_1 g_2) = \rho(g_1)\rho(g_2)$. A representation is **irreducible** if $\mathcal{H}$ has no nontrivial subspace that is invariant under all $\rho(g)$.

Schur's lemma. An operator that commutes with every element of an irreducible representation is a scalar multiple of the identity.

Corollary. If a Hamiltonian H commutes with every element of a representation of G, then H is a scalar on every irreducible subrepresentation. All states in a given irreducible block have the same energy — the symmetry enforces degeneracy.

Conservation as exact identity

Let $[G, H] = 0$. Then for any state $|\psi\rangle$ and time t, using $|\psi(t)\rangle = e^{-iHt/\hbar}|\psi\rangle$ and the commutator vanishing:

$$\langle\psi(t)|G|\psi(t)\rangle = \langle\psi|e^{iHt/\hbar}Ge^{-iHt/\hbar}|\psi\rangle = \langle\psi|G|\psi\rangle.$$

More generally, for any function f,

$$\langle\psi(t)|f(G)|\psi(t)\rangle = \langle\psi|f(G)|\psi\rangle,$$

so the entire probability distribution of G is conserved. Contrast this with the classical case, where Noether's theorem conserves only the expectation value (or the value along a trajectory).

Heisenberg picture statement

In the Heisenberg picture (Node 2.4),

$$\frac{dG_H}{dt} = \frac{i}{\hbar}[H, G_H].$$

The right-hand side vanishes if and only if $[H, G] = 0$. Symmetry and conservation are therefore the same statement in two notations.

Anti-unitary case — time reversal

The only exception to unitarity is time reversal, which is implemented by an anti-unitary operator T satisfying $THT^{-1} = H$ for a time-reversal-invariant Hamiltonian. Anti-unitarity flips the sign of imaginary parts (in particular, T conjugates complex numbers), which is the algebraic reflection of time's arrow.

Worked Example

SU(2) symmetry and angular momentum

A system of a spin-$\frac{1}{2}$ particle in a spherically symmetric potential has three conserved observables — the components of total angular momentum $\vec{J} = (J_x, J_y, J_z)$, generators of the $SU(2)$ rotation group. They satisfy

$$[J_a, J_b] = i\hbar\epsilon_{abc}J_c.$$

Spherical symmetry of the Hamiltonian means $[J_a, H] = 0$ for each component.

Good quantum numbers

$J^2 = J_x^2 + J_y^2 + J_z^2$ commutes with all J_a (Casimir invariant) and with H. A complete set of commuting observables is $\{H, J^2, J_z\}$, and energy eigenstates can be labelled $|n, j, m\rangle$ with $J^2 = \hbar^2 j(j+1)$ and $J_z = \hbar m$.

Degeneracy

Each energy level $E_{n,j}$ is $(2j+1)$-fold degenerate — one state for each value of $m \in \{-j, -j+1, \ldots, j\}$. This is not a coincidence but an exact consequence of $SU(2)$ symmetry via Schur's lemma: the $2j+1$ states form an irreducible representation of $SU(2)$, and the Hamiltonian must be a constant on this irreducible subspace.

Selection rules

An electric dipole perturbation $\vec{d} \cdot \vec{E}$ is a vector operator (spin-1 tensor). The Wigner–Eckart theorem then forces matrix elements between angular-momentum states to vanish unless $\Delta j \in \{-1, 0, +1\}$ and $\Delta m \in \{-1, 0, +1\}$. Forbidden transitions in atomic spectra are exactly those that would violate this symmetry-based selection rule.

Qubit example

For a single qubit, the symmetry group that commutes with $H = (\hbar\omega/2)Z$ is the $U(1)$ generated by Z itself. Since $[Z, Z] = 0$, the operator Z is conserved. This matches Larmor precession (Node 2.4): the Bloch vector precesses around the z-axis, leaving the z-component unchanged. The full three-dimensional $SU(2)$ symmetry is broken by the field; only its $U(1)$ subgroup survives.

Cross-Links

- **Linear Algebra:** Commuting operators are simultaneously diagonalizable; invariant subspaces under a set of operators.
- **Math-Intuition:** Groups, homomorphisms, and Lie algebras — the structural objects underlying symmetry.

- **Systems Thinking:** Conservation laws are system invariants; every "what stays fixed under the dynamics?" question in systems analysis has a quantum version as "what commutes with H?"
- **Quantum Thermodynamics (Module 9):** Conserved quantities define generalized Gibbs ensembles; chemical potentials are Lagrange multipliers for conserved charges.
- **Module 10 (Foundations):** Spontaneous symmetry breaking and the structural asymmetries of the Standard Model use this postulate as their starting point.
- **VQA:** Symmetry-preserving ansätze restrict to subspaces of the Hilbert space and give better optimization landscapes — a direct practical application of Schur's lemma.

Structural Compression

Invariant: A unitary symmetry commutes with the Hamiltonian if and only if its generator is exactly conserved.

Mechanism: $[G, H] = 0$ via the Heisenberg equation $\dot{G}_H = (i/\hbar)[H, G_H]$.

Cross-System Echo: Classical Noether theorem conserves the value of a Noether charge along trajectories. Quantum Noether theorem conserves the entire probability distribution of the corresponding operator — a strictly stronger statement, reflecting the algebraic rather than trajectory-based nature of quantum dynamics.

Exercises

1. Prove that if $[G, H] = 0$, then the expectation value $\langle\psi(t)|G|\psi(t)\rangle$ is independent of t.
2. Strengthen the previous result: show that if $[G, H] = 0$, then the full distribution $\{p(g_i)\}$ of outcomes of a G-measurement is conserved.
3. Verify the angular momentum commutation relations $[J_a, J_b] = i\hbar\epsilon_{abc}J_c$ for the spin-$\frac{1}{2}$ representation $J_a = (\hbar/2)\sigma_a$.
4. Show that J^2 commutes with each J_a. (Hint: use the commutation relations and expand.)
5. Apply Schur's lemma: if a Hermitian operator commutes with the 2×2 Pauli-generated $SU(2)$ representation on $\mathbb{C}^2$, show that it must be proportional to the identity.
6. Identify a continuous symmetry of the Hamiltonian $H = (\hbar\omega/2)Z$ on $\mathbb{C}^2$ and find its generator. How many dimensions does the symmetry group have?
7. Explain, structurally, why degeneracies in the spectrum of H always signal either a symmetry of H or an "accidental" coincidence, and give an example of each in two-level and three-level systems.