Chapter 2

Foundations I: States and Ensembles

2.1 Axioms of quantum mechanics

For a few lectures I have been talking about quantum this and that, but I have never defined what quantum theory is. It is time to correct that omission.

Quantum theory is a mathematical model of the physical world. To characterize the model, we need to specify how it will represent: states, observables, measurements, dynamics.

1. **States**. A state is a complete description of a physical system. In quantum mechanics, a state is a *ray* in a *Hilbert space.*

   What is a Hilbert space?

   a) It is a vector space over the complex numbers $\mathbf{C}$. Vectors will be denoted $|\psi\rangle$ (Dirac's ket notation).

   b) It has an inner product $\langle\psi|\varphi\rangle$ that maps an ordered pair of vectors to $\mathbf{C}$, defined by the properties

   (i) Positivity: $\langle\psi|\psi\rangle > 0$ for $|\psi\rangle = 0$

   (ii) Linearity: $\langle\varphi|(a|\psi_1\rangle + b|\psi_2\rangle) = a\langle\varphi|\psi_1\rangle + b\langle\varphi|\psi_2\rangle$

   (iii) Skew symmetry: $\langle\varphi|\psi\rangle = \langle\psi|\varphi\rangle^*$

   c) It is *complete* in the norm $||\psi|| = \langle\psi|\psi\rangle^{1/2}$

   (Completeness is an important proviso in infinite-dimensional function spaces, since it will ensure the convergence of certain eigenfunction expansions – e.g., Fourier analysis. But mostly we'll be content to work with finite-dimensional inner product spaces.)

   What is a ray? It is an equivalence class of vectors that differ by multiplication by a nonzero complex scalar. We can choose a representative of this class (for any nonvanishing vector) to have unit norm

$$\langle\psi|\psi\rangle = 1. \tag{2.1}$$

   We will also say that $|\psi\rangle$ and $e^{i\alpha}|\psi\rangle$ describe the same physical state, where $|e^{i\alpha}| = 1$.

   (Note that every ray corresponds to a possible state, so that given two states $|\varphi\rangle, |\psi\rangle$, we can form another as $a|\varphi\rangle + b|\psi\rangle$ (the "superposition principle"). The *relative* phase in this superposition *is* physically significant; we identify $a|\varphi\rangle + b|\varphi\rangle$ with $e^{i\alpha}(a|\varphi\rangle + b|\psi\rangle)$ but *not* with $a|\varphi\rangle + e^{i\alpha}b|\psi\rangle$.)

2. **Observables**. An observable is a property of a physical system that in principle can be measured. In quantum mechanics, an observable is a *self-adjoint operator*. An operator is a linear map taking vectors to vectors

$$\mathbf{A} : |\psi\rangle \to \mathbf{A}|\psi\rangle, \mathbf{A}\left(a|\psi\rangle + b|\psi\rangle\right) = a\mathbf{A}|\psi\rangle + b\mathbf{B}|\psi\rangle. \tag{2.2}$$

   The adjoint of the operator $\mathbf{A}$ is defined by

$$\langle\varphi|\mathbf{A}\psi\rangle = \langle\mathbf{A}^\dagger\varphi|\psi\rangle, \tag{2.3}$$

   for all vectors $|\varphi\rangle, |\psi\rangle$ (where here I have denoted $\mathbf{A}|\psi\rangle$ as $|\mathbf{A}\psi\rangle$). $\mathbf{A}$ is self-adjoint if $\mathbf{A} = \mathbf{A}^\dagger$.

   If $\mathbf{A}$ and $\mathbf{B}$ are self adjoint, then so is $\mathbf{A} + \mathbf{B}$ (because $(\mathbf{A} + \mathbf{B})^\dagger = \mathbf{A}^\dagger + \mathbf{B}^\dagger$) but $(\mathbf{AB})^\dagger = \mathbf{B}^\dagger\mathbf{A}^\dagger$, so $\mathbf{AB}$ is self adjoint only if $\mathbf{A}$ and $\mathbf{B}$ commute. Note that $\mathbf{AB}+\mathbf{BA}$ and $i(\mathbf{AB}-\mathbf{BA})$ are always self-adjoint if $\mathbf{A}$ and $\mathbf{B}$ are.

   A self-adjoint operator in a Hilbert space $\mathcal{H}$ has a spectral representation – it's eigenstates form a complete orthonormal basis in $\mathcal{H}$. We can express a self-adjoint operator $\mathbf{A}$ as

$$\mathbf{A} = \sum_n a_n \mathbf{P}_n. \tag{2.4}$$

Here each a_n is an eigenvalue of $\mathbf{A}$, and $\mathbf{P}_n$ is the corresponding orthogonal projection onto the space of eigenvectors with eigenvalue a_n. (If a_n is nondegenerate, then $\mathbf{P}_n = |n\rangle\langle n|$; it is the projection onto the corresponding eigenvector.) The $\mathbf{P}_n$'s satisfy

$$\begin{aligned} \mathbf{P}_n \mathbf{P}_m &= \delta_{n,m} \mathbf{P}_n \\ \mathbf{P}_n^\dagger &= \mathbf{P}_n. \end{aligned} \tag{2.5}$$

(For unbounded operators in an infinite-dimensional space, the definition of self-adjoint and the statement of the spectral theorem are more subtle, but this need not concern us.)

3. **Measurement**. In quantum mechanics, the numerical outcome of a measurement of the observable $\mathbf{A}$ is an eigenvalue of $\mathbf{A}$; right after the measurement, the quantum state is an eigenstate of $\mathbf{A}$ with the measured eigenvalue. If the quantum state just prior to the measurement is $|\psi\rangle$, then the outcome a_n is obtained with *probability*

$$\textbf{Prob}\ (a_n) = \parallel \mathbf{P}_n|\psi\rangle \parallel^2 = \langle\psi|\mathbf{P}_n|\psi\rangle; \tag{2.6}$$

If the outcome is a_n is attained, then the (normalized) quantum state becomes

$$\frac{\mathbf{P}_n|\psi\rangle}{\left(\langle\psi|\mathbf{P}_n|\psi\rangle\right)^{1/2}}. \tag{2.7}$$

(Note that if the measurement is immediately repeated, then according to this rule the same outcome is attained again, with probability one.)

4. **Dynamics**. Time evolution of a quantum state is unitary; it is generated by a self-adjoint operator, called the *Hamiltonian* of the system. In the *Schrödinger picture* of dynamics, the vector describing the system moves in time as governed by the *Schrödinger equation*

$$\frac{d}{dt}|\psi(t)\rangle = -i\mathbf{H}|\psi(t)\rangle, \tag{2.8}$$

where $\mathbf{H}$ is the Hamiltonian. We may reexpress this equation, to first order in the infinitesimal quantity dt, as

$$|\psi(t+dt)\rangle = (\mathbf{1} - i\mathbf{H}dt)|\psi(t)\rangle. \tag{2.9}$$

The operator $\mathbf{U}(dt) \equiv \mathbf{1} - i\mathbf{H}dt$ is unitary; because $\mathbf{H}$ is self-adjoint it satisfies $\mathbf{U}^\dagger\mathbf{U} = 1$ to linear order in dt. Since a product of unitary operators is finite, time evolution over a finite interval is also unitary

$$|\psi(t)\rangle = \mathbf{U}(t)|\psi(0)\rangle. \tag{2.10}$$

In the case where $\mathbf{H}$ is t-independent; we may write $\mathbf{U} = e^{-it\mathbf{H}}$.

This completes the mathematical formulation of quantum mechanics. We immediately notice some curious features. One oddity is that the Schrödinger equation is linear, while we are accustomed to nonlinear dynamical equations in classical physics. This property seems to beg for an explanation. But far more curious is the mysterious dualism; there are two quite distinct ways for a quantum state to change. On the one hand there is unitary evolution, which is deterministic. If we specify $|\psi(0)\rangle$, the theory predicts the state $|\psi(t)\rangle$ at a later time.

But on the other hand there is measurement, which is probabilistic. The theory does not make definite predictions about the measurement outcomes; it only assigns probabilities to the various alternatives. This is troubling, because it is unclear why the measurement process should be governed by different physical laws than other processes.

Beginning students of quantum mechanics, when first exposed to these rules, are often told not to ask "why?" There is much wisdom in this advice. But I believe that it can be useful to ask way. In future lectures. we will return to this disconcerting dualism between unitary evolution and measurement, and will seek a resolution.

2.2 The Qubit

The indivisible unit of classical information is the *bit*, which takes one of the two possible values $\{0,1\}$. The corresponding unit of quantum information is called the "quantum bit" or *qubit*. It describes a state in the simplest possible quantum system.

The smallest nontrivial Hilbert space is two-dimensional. We may denote an orthonormal basis for a two-dimensional vector space as $\{|0\rangle, |1\rangle\}$. Then the most general normalized state can be expressed as

$$a|0\rangle + b|1\rangle, \tag{2.11}$$

where a, b are complex numbers that satisfy $|a|^2 + |b|^2 = 1$, and the overall phase is physically irrelevant. A *qubit* is a state in a two-dimensional Hilbert space that can take any value of the form eq. (2.11).

We can perform a measurement that projects the qubit onto the basis $\{|0\rangle, |1\rangle\}$. Then we will obtain the outcome $|0\rangle$ with probability $|a|^2$, and the outcome $|1\rangle$ with probability $|b|^2$. Furthermore, except in the cases $a = 0$ and $b = 0$, the measurement irrevocably disturbs the state. If the value of the qubit is initially unknown, then there is no way to determine a and b with that single measurement, or any other conceivable measurement. However, *after* the measurement, the qubit has been prepared in a *known* state – either $|0\rangle$ or $|1\rangle$ – that differs (in general) from its previous state.

In this respect, a qubit differs from a classical bit; we can measure a classical bit without disturbing it, and we can decipher all of the information that it encodes. But suppose we have a classical bit that really does have a definite value (either 0 or 1), but that value is initially unknown to us. Based on the information available to us we can only say that there is a *probability* p_0 that the bit has the value 0, and a probability p_1 that the bit has the value 1, where $p_0 + p_1 = 1$. When we measure the bit, we acquire additional information; afterwards we know the value with 100% confidence.

An important question is: what is the essential difference between a qubit and a *probabilistic* classical bit? In fact they are *not* the same, for several reasons that we will explore.

2.2.1 Spin-$\frac{1}{2}$

First of all, the coefficients a and b in eq. (2.11) encode more than just the probabilities of the outcomes of a measurement in the $\{|0\rangle, |1\rangle\}$ basis. In particular, the *relative phase* of a and b also has physical significance.

For a physicist, it is natural to interpret eq. (2.11) as the spin state of an object with spin-$\frac{1}{2}$ (like an electron). Then $|0\rangle$ and $|1\rangle$ are the spin up ($|\uparrow\rangle$) and spin down ($|\downarrow\rangle$) states along a particular axis such as the z-axis. The two real numbers characterizing the qubit (the complex numbers a and b, modulo the normalization and overall phase) describe the *orientation* of the spin in three-dimensional space (the polar angle θ and the azimuthal angle φ).

We cannot go deeply here into the theory of symmetry in quantum mechanics, but we will briefly recall some elements of the theory that will prove useful to us. A symmetry is a transformation that acts on a state of a system,

yet leaves all observable properties of the system unchanged. In quantum mechanics, observations are measurements of self-adjoint operators. If $\mathbf{A}$ is measured in the state $|\psi\rangle$, then the outcome $|a\rangle$ (an eigenvector of $\mathbf{A}$) occurs with probability $|\langle a|\psi\rangle|^2$. A symmetry should leave these probabilities unchanged (when we "rotate" both the system *and* the apparatus).

A symmetry, then, is a mapping of vectors in Hilbert space

$$|\psi\rangle \to |\psi'\rangle, \tag{2.12}$$

that preserves the absolute values of inner products

$$|\langle\varphi|\psi\rangle| = |\langle\varphi'|\psi'\rangle|, \tag{2.13}$$

for all $|\varphi\rangle$ and $|\psi\rangle$. According to a famous theorem due to Wigner, a mapping with this property can always be chosen (by adopting suitable phase conventions) to be either unitary or antiunitary. The antiunitary alternative, while important for discrete symmetries, can be excluded for continuous symmetries. Then the symmetry acts as

$$|\psi\rangle \to |\psi'\rangle = \mathbf{U}|\psi\rangle, \tag{2.14}$$

where $\mathbf{U}$ is unitary (and in particular, *linear*).

Symmetries form a group: a symmetry transformation can be inverted, and the product of two symmetries is a symmetry. For each symmetry operation R acting on our physical system, there is a corresponding unitary transformation $\mathbf{U}(R)$. Multiplication of these unitary operators must respect the group multiplication law of the symmetries – applying $R_1 \circ R_2$ should be equivalent to first applying R_2 and subsequently R_1. Thus we demand

$$\mathbf{U}(R_1)\mathbf{U}(R_2) = \text{ Phase } (R_1, R_2)\mathbf{U}(R_1 \circ R_2) \tag{2.15}$$

The phase is permitted in eq. (2.15) because quantum states are *rays*; we need only demand that $\mathbf{U}(R_1 \circ R_2)$ act the same way as $\mathbf{U}(R_1)\mathbf{U}(R_2)$ on rays, not on vectors. $\mathbf{U}(R)$ provides a unitary representation (up to a phase) of the symmetry group.

So far, our concept of symmetry has no connection with dynamics. Usually, we demand of a symmetry that it respect the dynamical evolution of the system. This means that it should not matter whether we first transform the system and then evolve it, or first evolve it and then transform it. In other words, the diagram

$$\begin{tikzcd}
\boxed{\text{Initial}} \arrow[r, "dynamics"] \arrow[d, "rotation"] & \boxed{\text{Final}} \arrow[d, "rotation"] \\
\boxed{\text{New Initial}} \arrow[r, "dynamics"] & \boxed{\text{New Final}}
\end{tikzcd}$$

is commutative. This means that the time evolution operator $e^{it\mathbf{H}}$ should commute with the symmetry transformation $\mathbf{U}(R)$:

$$\mathbf{U}(R)e^{-it\mathbf{H}} = e^{-it\mathbf{H}}\mathbf{U}(R), \tag{2.16}$$

and expanding to linear order in t we obtain

$$\mathbf{U}(R)\mathbf{H} = \mathbf{H}\mathbf{U}(R) \tag{2.17}$$

For a continuous symmetry, we can choose R infinitesimally close to the identity, $R = I + \epsilon T$, and then $\mathbf{U}$ is close to $\mathbf{1}$,

$$\mathbf{U} = \mathbf{1} - i\varepsilon\mathbf{Q} + O(\varepsilon^2). \tag{2.18}$$

From the unitarity of $\mathbf{U}$ (to order ε) it follows that $\mathbf{Q}$ is an observable, $\mathbf{Q} = \mathbf{Q}^\dagger$. Expanding eq. (2.17) to linear order in ε we find

$$[\mathbf{Q}, \mathbf{H}] = 0; \tag{2.19}$$

the observable $\mathbf{Q}$ commutes with the Hamiltonian.

Eq. (2.19) is a *conservation law*. It says, for example, that if we prepare an eigenstate of $\mathbf{Q}$, then time evolution governed by the Schrödinger equation will preserve the eigenstate. We have seen that symmetries imply conservation laws. Conversely, given a conserved quantity $\mathbf{Q}$ satisfying eq. (2.19) we can construct the corresponding symmetry transformations. Finite transformations can be built as a product of many infinitesimal ones

$$R = (1 + \frac{\theta}{N}T)^N \Rightarrow \mathbf{U}(R) = (\mathbf{1} + i\frac{\theta}{N}\mathbf{Q})^N \to e^{i\theta\mathbf{Q}}, \tag{2.20}$$

(taking the limit $N \to \infty$). Once we have decided how infinitesimal symmetry transformations are represented by unitary operators, then it is also

determined how finite transformations are represented, for these can be built as a product of infinitesimal transformations. We say that $\mathbf{Q}$ is the *generator* of the symmetry.

Let us briefly recall how this general theory applies to spatial rotations and angular momentum. An infinitesimal rotation by $d\theta$ about the axis specified by the unit vector $\hat{n} = (n_1, n_2, n_3)$ can be expressed as

$$R(\hat{n}, d\theta) = I - id\theta\hat{n}\cdot\vec{J}, \tag{2.21}$$

where (J_1, J_2, J_3) are the components of the angular momentum. A finite rotation is expressed as

$$R(\hat{n}, \theta) = \exp(-i\theta\hat{n}\cdot\vec{J}). \tag{2.22}$$

Rotations about distinct axes don't commute. From elementary properties of rotations, we find the commutation relations

$$[J_k, J_\ell] = i\varepsilon_{k\ell m}J_m, \tag{2.23}$$

where $\varepsilon_{k\ell m}$ is the totally antisymmetric tensor with $\varepsilon_{123} = 1$, and repeated indices are summed. To implement rotations on a quantum system, we find self-adjoint operators $\mathbf{J}_1, \mathbf{J}_2, \mathbf{J}_3$ in Hilbert space that satisfy these relations.

The "defining" representation of the rotation group is three dimensional, but the simplest nontrivial irreducible representation is two dimensional, given by

$$\mathbf{J}_k = \frac{1}{2}\boldsymbol{\sigma}_k, \tag{2.24}$$

where

$$\boldsymbol{\sigma}_1 = \begin{pmatrix} 0 & 1 \\ 1 & 0 \end{pmatrix}, \boldsymbol{\sigma}_2 = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix}, \boldsymbol{\sigma}_3 = \begin{pmatrix} 1 & 0 \\ 0 & -1 \end{pmatrix}, \tag{2.25}$$

are the Pauli matrices. This is the unique two-dimensional irreducible representation, up to a unitary change of basis. Since the eigenvalues of $\mathbf{J}_k$ are $\pm\frac{1}{2}$, we call this the spin-$\frac{1}{2}$ representation. (By identifying $\mathbf{J}$ as the angular-momentum, we have implicitly chosen units with $\hbar = 1$).

The Pauli matrices also have the properties of being mutually anticommuting and squaring to the identity,

$$\boldsymbol{\sigma}_k\boldsymbol{\sigma}_\ell + \boldsymbol{\sigma}_\ell\boldsymbol{\sigma}_k = 2\delta_{k\ell}\mathbf{1}, \tag{2.26}$$

So we see that $(\hat{n} \cdot \vec{\boldsymbol{\sigma}})^2 = n_k n_\ell \boldsymbol{\sigma}_k \boldsymbol{\sigma}_\ell = n_k n_k \mathbf{1} = \mathbf{1}$. By expanding the exponential series, we see that finite rotations are represented as

$$\mathbf{U}(\hat{n}, \theta) = e^{-i\frac{\theta}{2}\hat{n}\cdot\vec{\boldsymbol{\sigma}}} = \mathbf{1}\cos\frac{\theta}{2} - i\hat{n}\cdot\vec{\boldsymbol{\sigma}}\sin\frac{\theta}{2}. \tag{2.27}$$

The most general 2×2 unitary matrix with determinant 1 can be expressed in this form. Thus, we are entitled to think of a qubit as the state of a spin-$\frac{1}{2}$ object, and an arbitrary unitary transformation acting on the state (aside from a possible rotation of the overall phase) is a *rotation* of the spin.

A peculiar property of the representation $\mathbf{U}(\hat{n}, \theta)$ is that it is *double-valued*. In particular a rotation by 2π about any axis is represented nontrivially:

$$\mathbf{U}(\hat{n}, \theta = 2\pi) = -\mathbf{1}. \tag{2.28}$$

Our representation of the rotation group is really a representation "up to a sign"

$$\mathbf{U}(R_1)\mathbf{U}(R_2) = \pm\mathbf{U}(R_1 \circ R_2). \tag{2.29}$$

But as already noted, this is acceptable, because the group multiplication is respected on *rays*, though not on vectors. These double-valued representations of the rotation group are called *spinor* representations. (The existence of spinors follows from a topological property of the group — it is not simply connected.)

While it is true that a rotation by 2π has no detectable effect on a spin-$\frac{1}{2}$ object, it would be wrong to conclude that the spinor property has no observable consequences. Suppose I have a machine that acts on a pair of spins. If the first spin is up, it does nothing, but if the first spin is down, it rotates the second spin by 2π. Now let the machine act when the first spin is in a *superposition* of up and down. Then

$$\frac{1}{\sqrt{2}}\left(|\uparrow\rangle_1 + |\downarrow\rangle_1\right)|\uparrow\rangle_2 \to \frac{1}{\sqrt{2}}\left(|\uparrow\rangle_1 - |\downarrow\rangle_1\right)|\uparrow\rangle_2 \,. \tag{2.30}$$

While there is no detectable effect on the second spin, the state of the first has flipped to an orthogonal state, which is very much observable.

In a rotated frame of reference, a rotation $R(\hat{n}, \theta)$ becomes a rotation through the same angle but about a rotated axis. It follows that the three components of angular momentum transform under rotations as a vector:

$$\mathbf{U}(R)\mathbf{J}_k\mathbf{U}(R)^\dagger = R_{k\ell}\mathbf{J}_\ell. \tag{2.31}$$

Thus, if a state $|m\rangle$ is an eigenstate of $\mathbf{J}_3$

$$\mathbf{J}_3|m\rangle = m|m\rangle, \tag{2.32}$$

then $\mathbf{U}(R)|m\rangle$ is an eigenstate of $R\mathbf{J}_3$ with the same eigenvalue:

$$\begin{aligned} R\mathbf{J}_3\left(\mathbf{U}(R)|m\rangle\right) &= \mathbf{U}(R)\mathbf{J}_3\mathbf{U}(R)^\dagger\mathbf{U}(R)|m\rangle \\ &= \mathbf{U}(R)\mathbf{J}_3|m\rangle = m\left(\mathbf{U}(R)|m\rangle\right). \end{aligned} \tag{2.33}$$

Therefore, we can construct eigenstates of angular momentum along the axis $\hat{n} = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$ by applying a rotation through θ, about the axis $\hat{n}' = (-\sin\varphi, \cos\varphi, 0)$, to a $\mathbf{J}_3$ eigenstate. For our spin-$\frac{1}{2}$ representation, this rotation is

$$\begin{aligned} \exp\left[-i\frac{\theta}{2}\hat{n}'\cdot\vec{\boldsymbol{\sigma}}\right] &= \exp\left[\frac{\theta}{2}\begin{pmatrix} 0 & -e^{-i\varphi} \\ e^{i\varphi} & 0 \end{pmatrix}\right] \\ &= \begin{pmatrix} \cos\frac{\theta}{2} & -e^{-i\varphi}\sin\frac{\theta}{2} \\ e^{i\varphi}\sin\frac{\theta}{2} & \cos\frac{\theta}{2} \end{pmatrix}, \end{aligned} \tag{2.34}$$

and applying it to $\binom{1}{0}$, the $\mathbf{J}_3$ eigenstate with eigenvalue 1, we obtain

$$|\psi(\theta,\varphi)\rangle = \begin{pmatrix} e^{-i\varphi/2}\cos\frac{\theta}{2} \\ e^{i\varphi/2}\sin\frac{\theta}{2} \end{pmatrix}, \tag{2.35}$$

(up to an overall phase). We can check directly that this is an eigenstate of

$$\hat{n}\cdot\vec{\boldsymbol{\sigma}} = \begin{pmatrix} \cos\theta & e^{-i\varphi}\sin\theta \\ e^{i\varphi}\sin\theta & -\cos\theta \end{pmatrix}, \tag{2.36}$$

with eigenvalue one. So we have seen that eq. (2.11) with $a = e^{-i\varphi/2}\cos\frac{\theta}{2}$, $b = e^{i\varphi/2}\sin\frac{\theta}{2}$, can be interpreted as a spin pointing in the (θ, φ) direction.

We noted that we cannot determine a and b with a single measurement. Furthermore, even with many identical copies of the state, we cannot completely determine the state by measuring each copy only along the z-axis. This would enable us to estimate $|a|$ and $|b|$, but we would learn nothing about the relative phase of a and b. Equivalently, we would find the component of the spin along the z-axis

$$\langle\psi(\theta,\varphi)|\sigma_3|\psi(\theta,\varphi)\rangle = \cos^2\frac{\theta}{2} - \sin^2\frac{\theta}{2} = \cos\theta, \tag{2.37}$$

but we would not learn about the component in the $x-y$ plane. The problem of determining $|\psi\rangle$ by measuring the spin is equivalent to determining the unit vector $\hat{n}$ by measuring its components along various axes. Altogether, measurements along three different axes are required. *E.g.*, from $\langle\sigma_3\rangle$ and $\langle\sigma_1\rangle$ we can determine n_3 and n_1, but the sign of n_2 remains undetermined. Measuring $\langle\sigma_2\rangle$ would remove this remaining ambiguity.

Of course, if we are permitted to rotate the spin, then only measurements along the z-axis will suffice. That is, measuring a spin along the $\hat{n}$ axis is equivalent to first applying a rotation that rotates the $\hat{n}$ axis to the axis $\hat{z}$, and then measuring along $\hat{z}$.

In the special case $\theta = \frac{\pi}{2}$ and $\varphi = 0$ (the $\hat{x}$-axis) our spin state is

$$|\uparrow_x\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow_z\rangle + |\downarrow_z\rangle\right), \tag{2.38}$$

("spin-up along the x-axis"). The orthogonal state ("spin down along the x-axis") is

$$|\downarrow_x\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow_z\rangle - |\downarrow_z\rangle\right). \tag{2.39}$$

For either of these states, if we measure the spin along the z-axis, we will obtain $|\uparrow_z\rangle$ with probability $\frac{1}{2}$ and $|\downarrow_z\rangle$ with probability $\frac{1}{2}$.

Now consider the combination

$$\frac{1}{\sqrt{2}}\left(|\uparrow_x\rangle + |\downarrow_x\rangle\right). \tag{2.40}$$

This state has the property that, if we measure the spin along the x-axis, we obtain $|\uparrow_x\rangle$ or $|\downarrow_x\rangle$, each with probability $\frac{1}{2}$. Now we may ask, what if we measure the state in eq. (2.40) along the z-axis?

If these were probabilistic classical bits, the answer would be obvious. The state in eq. (2.40) is in one of two states, and for *each* of the two, the probability is $\frac{1}{2}$ for pointing up or down along the z-axis. So of course we should find up with probability $\frac{1}{2}$ when we measure along the z-axis.

But not so for qubits! By adding eq. (2.38) and eq. (2.39), we see that the state in eq. (2.40) is really $|\uparrow_z\rangle$ in disguise. When we measure along the z-axis, we always find $|\uparrow_z\rangle$, never $|\downarrow_z\rangle$.

We see that for qubits, as opposed to probabilistic classical bits, probabilities can add in unexpected ways. This is, in its simplest guise, the phenomenon called "quantum interference," an important feature of quantum information.

It should be emphasized that, while this *formal* equivalence with a spin-$\frac{1}{2}$ object applies to any two-level quantum system, of course not every two-level system transforms as a spinor under rotations!

2.2.2 Photon polarizations

Another important two-state system is provided by a *photon*, which can have two independent polarizations. These photon polarization states also transform under rotations, but photons differ from our spin-$\frac{1}{2}$ objects in two important ways: (1) Photons are massless. (2) Photons have spin-1 (they are not spinors).

Now is not a good time for a detailed discussion of the unitary representations of the Poincare group. Suffice it to say that the *spin* of a particle classifies how it transforms under the *little group*, the subgroup of the Lorentz group that preserves the particle's momentum. For a massive particle, we may always boost to the particle's rest frame, and then the little group is the rotation group.

For massless particles, there is no rest frame. The finite-dimensional unitary representations of the little group turn out to be representations of the rotation group in *two* dimensions, the rotations about the axis determined by the momentum. Of course, for a photon, this corresponds to the familiar property of classical light – the waves are polarized transverse to the direction of propagation.

Under a rotation about the axis of propagation, the two linear polarization states ($|x\rangle$ and $|y\rangle$ for horizontal and vertical polarization) transform as

$$\begin{aligned} |x\rangle &\to \cos\theta|x\rangle + \sin\theta|y\rangle \\ |y\rangle &\to -\sin\theta|x\rangle + \cos\theta|y\rangle. \end{aligned} \tag{2.41}$$

This two-dimensional representation is actually reducible. The matrix

$$\begin{pmatrix} \cos\theta & \sin\theta \\ -\sin\theta & \cos\theta \end{pmatrix} \tag{2.42}$$

has the eigenstates

$$|R\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} 1 \\ i \end{pmatrix} \qquad |L\rangle = \frac{1}{\sqrt{2}}\begin{pmatrix} i \\ 1 \end{pmatrix}, \tag{2.43}$$

with eigenvalues $e^{i\theta}$ and $e^{-i\theta}$, the states of right and left circular polarization. That is, these are the eigenstates of the rotation generator

$$J = \begin{pmatrix} 0 & -i \\ i & 0 \end{pmatrix} = \sigma_y, \tag{2.44}$$

with eigenvalues ± 1. Because the eigenvalues are ± 1 (*not* $\pm \frac{1}{2}$) we say that the photon has spin-1.

In this context, the quantum interference phenomenon can be described this way: Suppose that we have a polarization analyzer that allows only one of the two linear photon polarizations to pass through. Then an x or y polarized photon has prob $\frac{1}{2}$ of getting through a 45^o rotated polarizer, and a 45^o polarized photon has prob $\frac{1}{2}$ of getting through an x and y analyzer. But an x photon *never* passes through a y analyzer. If we put a 45^o rotated analyzer in between an x and y analyzer, then $\frac{1}{2}$ the photons make it through each analyzer. But if we remove the analyzer in the middle *no* photons make it through the y analyzer.

A device can be constructed easily that rotates the linear polarization of a photon, and so applies the transformation Eq. (2.41) to our qubit. As noted, this is not the most general possible unitary transformation. But if we also have a device that alters the relative phase of the two orthogonal linear polarization states

$$\begin{aligned} |x\rangle &\to e^{i\omega/2}|x\rangle \\ |y\rangle &\to e^{-i\omega/2}|y\rangle \ , \end{aligned} \tag{2.45}$$

the two devices can be employed together to apply an arbitrary 2×2 unitary transformation (of determinant 1) to the photon polarization state.

2.3 The density matrix

2.3.1 The bipartite quantum system

The last lecture was about one qubit. This lecture is about *two* qubits. (Guess what the next lecture will be about!) Stepping up from one qubit to two is a bigger leap than you might expect. Much that is weird and wonderful about quantum mechanics can be appreciated by considering the properties of the quantum states of two qubits.

The axioms of §2.1 provide a perfectly acceptable general formulation of the quantum theory. Yet under many circumstances, we find that the axioms appear to be violated. The trouble is that our axioms are intended to characterize the quantum behavior of the entire universe. Most of the time, we are not so ambitious as to attempt to understand the physics of the whole universe; we are content to observe just our little corner. In practice, then, the observations we make are always limited to a small part of a much larger quantum system.

In the next several lectures, we will see that, when we limit our attention to just part of a larger system, then (contrary to the axioms):

1. States are *not* rays.

2. Measurements are *not* orthogonal projections.

3. Evolution is *not* unitary.

We can best understand these points by considering the simplest possible example: a two-qubit world in which we observe only one of the qubits.

So consider a system of two qubits. Qubit A is here in the room with us, and we are free to observe or manipulate it any way we please. But qubit B is locked in a vault where we can't get access to it. Given some quantum state of the two qubits, we would like to find a compact way to characterize the observations that can be made on qubit A alone.

We'll use $\{|0\rangle_A, |1\rangle_A\}$ and $\{|0\rangle_B, |1\rangle_B\}$ to denote orthonormal bases for qubits A and B respectively. Consider a quantum state of the two-qubit world of the form

$$|\psi\rangle_{AB} = a|0\rangle_A \otimes |0\rangle_B + b|1\rangle_A \otimes |1\rangle_B. \tag{2.46}$$

In this state, qubits A and B are *correlated.* Suppose we measure qubit A by projecting onto the $\{|0\rangle_A, |1\rangle_A\}$ basis. Then with probability $|a|^2$ we obtain the result $|0\rangle_A$, and the measurement prepares the state

$$|0\rangle_A \otimes |0\rangle_B. \tag{2.47}$$

with probability $|b|^2$, we obtain the result $|1\rangle_A$ and prepare the state

$$|1\rangle_A \otimes |1\rangle_B. \tag{2.48}$$

In either case, a definite state of qubit B is picked out by the measurement. If we subsequently measure qubit B, then we are guaranteed (with probability one) to find $|0\rangle_B$ if we had found $|0\rangle_A$, and we are guaranteed to find $|1\rangle_B$ if we found $|1\rangle_A$. In this sense, the outcomes of the $\{|0\rangle_A, |1\rangle_A\}$ and $\{|0\rangle_B, |1\rangle_B\}$ measurements are perfectly correlated in the state $|\psi\rangle_{AB}$.

But now I would like to consider more general observables acting on qubit A, and I would like to characterize the measurement outcomes for A alone (irrespective of the outcomes of any measurements of the inaccessible qubit B). An observable acting on qubit A only can be expressed as

$$\mathbf{M}_A \otimes \mathbf{1}_B, \tag{2.49}$$

where $\mathbf{M}_A$ is a self-adjoint operator acting on A, and $\mathbf{1}_B$ is the identity operator acting on B. The expectation value of the observable in the state $|\psi\rangle$ is:

$$\begin{aligned} &\langle\psi|\mathbf{M}_A \otimes \mathbf{1}_B|\psi\rangle \\ &= (a^*{}_A\langle 0| \otimes {}_B\langle 0| + b^*{}_B\langle 1| \otimes {}_B\langle 1|)\,(\mathbf{M}_A \otimes \mathbf{1}_B) \\ &\quad (a|0\rangle_A \otimes |0\rangle_B + b|1\rangle_A \otimes |1\rangle_B) \\ &= |a|^2{}_A\langle 0|\mathbf{M}_A|0\rangle_A + |b|^2{}_A\langle 1|\mathbf{M}_A|1\rangle_A, \end{aligned} \tag{2.50}$$

(where we have used the orthogonality of $|0\rangle_B$ and $|1\rangle_B$). This expression can be rewritten in the form

$$\langle\mathbf{M}_A\rangle = \mathrm{tr}\,(\mathbf{M}_A\boldsymbol{\rho}_A)\,, \tag{2.51}$$

$$\boldsymbol{\rho}_A = |a|^2|0\rangle_A\,{}_A\langle 0| + |b|^2|1\rangle_A\,{}_A\langle 1|, \tag{2.52}$$

and $\mathrm{tr}(\cdot)$ denotes the *trace*. The operator $\boldsymbol{\rho}_A$ is called the *density operator* (or *density matrix*) for qubit A. It is self-adjoint, positive (its eigenvalues are nonnegative) and it has unit trace (because $|\psi\rangle$ is a normalized state.)

Because $\langle\mathbf{M}_A\rangle$ has the form eq. (2.51) for *any* observable $\mathbf{M}_A$ acting on qubit A, it is consistent to interpret $\boldsymbol{\rho}_A$ as representing an *ensemble* of possible quantum states, each occurring with a specified probability. That is, we would obtain precisely the same result for $\langle\mathbf{M}_A\rangle$ if we stipulated that qubit A is in one of two quantum states. With probability $p_0 = |a|^2$ it is in the quantum state $|0\rangle_A$, and with probability $p_1 = |b|^2$ it is in the state

$|1\rangle_A$. If we are interested in the result of any possible measurement, we can consider $\mathbf{M}_A$ to be the projection $\mathbf{E}_A(a)$ onto the relevant eigenspace of a particular observable. Then

$$\mathrm{Prob}\ (a) = p_0\ {}_A\langle 0|\mathbf{E}_A(a)|0\rangle_A + p_1\ {}_A\langle 1|\mathbf{E}_A(a)|1\rangle_A, \tag{2.53}$$

which is the probability of outcome a summed over the ensemble, and weighted by the probability of each state in the ensemble.

We have emphasized previously that there is an essential difference between a coherent superposition of the states $|0\rangle_A$ and $|1\rangle_A$, and a probabilistic ensemble, in which $|0\rangle_A$ and $|1\rangle_A$ can each occur with specified probabilities. For example, for a spin-$\frac{1}{2}$ object we have seen that if we measure $\boldsymbol{\sigma}_1$ in the state $\frac{1}{\sqrt{2}}(|\uparrow_z\rangle + |\downarrow_z\rangle)$, we will obtain the result $|\uparrow_x\rangle$ with probability one. But the ensemble in which $|\uparrow_z\rangle$ and $|\downarrow_z\rangle$ each occur with probability $\frac{1}{2}$ is represented by the density operator

$$\begin{aligned} \boldsymbol{\rho} &= \frac{1}{2}(|\uparrow_z\rangle\langle\uparrow_z| + |\downarrow_z\rangle\langle\downarrow_z|) \\ &= \frac{1}{2}\mathbf{1}, \end{aligned} \tag{2.54}$$

and the projection onto $|\uparrow_x\rangle$ then has the expectation value

$$\mathrm{tr}\,(|\uparrow_x\rangle\langle\uparrow_x|\boldsymbol{\rho}) = \frac{1}{2}. \tag{2.55}$$

In fact, we have seen that any state of one qubit represented by a ray can be interpreted as a spin pointing in some definite direction. But because the identity is left unchanged by any unitary change of basis, and the state $|\psi(\theta,\varphi)\rangle$ can be obtained by applying a suitable unitary transformation to $|\uparrow_z\rangle$, we see that for $\boldsymbol{\rho}$ given by eq. (2.54), we have

$$\mathrm{tr}\,(|\psi(\theta,\varphi)\rangle\langle\psi(\theta,\varphi)|\boldsymbol{\rho}) = \frac{1}{2}. \tag{2.56}$$

Therefore, if the state $|\psi\rangle_{AB}$ in eq. (2.57) is prepared, with $|a|^2 = |b|^2 = \frac{1}{2}$, and we measure the spin A along *any* axis, we obtain a completely random result; spin up or spin down can occur, each with probability $\frac{1}{2}$.

This discussion of the correlated two-qubit state $|\psi\rangle_{AB}$ is easily generalized to an arbitrary state of any bipartite quantum system (a system divided into two parts). The Hilbert space of a bipartite system is $\mathcal{H}_A \otimes \mathcal{H}_B$ where

$\mathcal{H}_{A,B}$ are the Hilbert spaces of the two parts. This means that if $\{|i\rangle_A\}$ is an orthonormal basis for $\mathcal{H}_A$ and $\{|\mu\rangle_B\}$ is an orthonormal basis for $\mathcal{H}_B$, then $\{|i\rangle_A \otimes |\mu\rangle_B\}$ is an orthonormal basis for $\mathcal{H}_A \otimes \mathcal{H}_B$. Thus an arbitrary pure state of $\mathcal{H}_A \otimes \mathcal{H}_B$ can be expanded as

$$|\psi\rangle_{AB} = \sum_{i,\mu} a_{i\mu} |i\rangle_A \otimes |\mu\rangle_B, \tag{2.57}$$

where $\sum_{i,\mu} |a_{i\mu}|^2 = 1$. The expectation value of an observable $\mathbf{M}_A \otimes \mathbf{1}_B$, that acts only on subsystem A is

$$\begin{aligned}\langle \mathbf{M}_A \rangle &= {}_{AB}\langle \psi | \mathbf{M}_A \otimes \mathbf{1}_B | \psi \rangle_{AB} \\ &= \sum_{j,\nu} a^*_{j\nu} \left({}_A\langle j| \otimes {}_B\langle \nu|\right) (\mathbf{M}_A \otimes \mathbf{1}_B) \sum_{i,\mu} a_{i\mu} \left(|i\rangle_A \otimes |\mu\rangle_B\right) \\ &= \sum_{i,j,\mu} a^*_{j\mu} a_{i\mu}\ {}_A\langle j | \mathbf{M}_A | i \rangle_A \\ &= \mathrm{tr}\left(\mathbf{M}_A \boldsymbol{\rho}_A\right), \end{aligned} \tag{2.58}$$

where

$$\begin{aligned}\boldsymbol{\rho}_A &= \mathrm{tr}_B \left(|\psi\rangle_{AB}\ {}_{AB}\langle \psi|\right) \\ &\equiv \sum_{i,j,\mu} a_{i\mu} a^*_{j\mu} |i\rangle_A\ {}_A\langle j| \ . \end{aligned} \tag{2.59}$$

We say that the density operator $\boldsymbol{\rho}_A$ for subsystem A is obtained by performing a partial *trace* over subsystem B of the density matrix (in this case a pure state) for the combined system AB.

From the definition eq. (2.59), we can immediately infer that $\boldsymbol{\rho}_A$ has the following properties:

1. $\boldsymbol{\rho}_A$ is self-adjoint: $\boldsymbol{\rho}_A = \boldsymbol{\rho}_A^\dagger$.
2. $\boldsymbol{\rho}_A$ is positive: For any $|\psi\rangle_A$ ${}_A\langle \psi | \boldsymbol{\rho}_A | \psi \rangle_A = \sum_\mu |\sum_i a_{i\mu}\ {}_A\langle \psi | i \rangle_A|^2 \geq 0$.
3. $\mathrm{tr}(\boldsymbol{\rho}_A) = 1$: We have tr $\boldsymbol{\rho}_A = \sum_{i,\mu} |a_{i\mu}|^2 = 1$, since $|\psi\rangle_{AB}$ is normalized.

It follows that $\boldsymbol{\rho}_A$ can be diagonalized, that the eigenvalues are all real and nonnegative, and that the eigenvalues sum to one.

If we are looking at a subsystem of a larger quantum system, then, even if the state of the larger system is a ray, the state of the subsystem need

not be; in general, the state is represented by a density operator. In the case where the state of the subsystem *is* a ray, and we say that the state is *pure*. Otherwise the state is *mixed*. If the state is a pure state $|\psi\rangle_A$, then the density matrix $\boldsymbol{\rho}_A = |\psi\rangle_A\ {}_A\langle\psi|$ is the *projection* onto the one-dimensional space spanned by $|\psi\rangle_A$. Hence a pure density matrix has the property $\boldsymbol{\rho}^2 = \boldsymbol{\rho}$. A general density matrix, expressed in the basis in which it is diagonal, has the form

$$\boldsymbol{\rho}_A = \sum_a p_a |\psi_a\rangle\langle\psi_a|, \tag{2.60}$$

where $0 < p_a \leq 1$ and $\sum_a p_a = 1$. If the state is not pure, there are two or more terms in this sum, and $\boldsymbol{\rho}^2 \neq \boldsymbol{\rho}$; in fact, tr $\boldsymbol{\rho}^2 = \sum p_a^2 < \sum p_a = 1$. We say that $\boldsymbol{\rho}$ is an *incoherent* superposition of the states $\{|\psi_a\rangle\}$; incoherent meaning that the relative phases of the $|\psi_a\rangle$ are experimentally inaccessible.

Since the expectation value of *any* observable $\mathbf{M}$ acting on the subsystem can be expressed as

$$\langle \mathbf{M} \rangle = \mathrm{tr}\mathbf{M}\boldsymbol{\rho} = \sum_a p_a \langle \psi_a | \mathbf{M} | \psi_a \rangle, \tag{2.61}$$

we see as before that we may interpret $\boldsymbol{\rho}$ as describing an *ensemble* of pure quantum states, in which the state $|\psi_a\rangle$ occurs with probability p_a. We have, therefore, come a long part of the way to understanding how probabilities arise in quantum mechanics when a quantum system A interacts with another system B. A and B become *entangled*, that is, correlated. The entanglement *destroys the coherence* of a superposition of states of A, so that some of the phases in the superposition become inaccessible if we look at A alone. We may describe this situation by saying that the state of system A *collapses* — it is in one of a set of alternative states, each of which can be assigned a probability.

2.3.2 Bloch sphere

Let's return to the case in which system A is a single qubit, and consider the form of the general density matrix. The most general self-adjoint 2×2 matrix has four real parameters, and can be expanded in the basis $\{\mathbf{1}, \boldsymbol{\sigma}_1, \boldsymbol{\sigma}_2, \boldsymbol{\sigma}_3\}$. Since each $\boldsymbol{\sigma}_i$ is traceless, the coefficient of $\mathbf{1}$ in the expansion of a density matrix $\boldsymbol{\rho}$ must be $\frac{1}{2}$ (so that $\mathrm{tr}(\boldsymbol{\rho}) = 1$), and $\boldsymbol{\rho}$ may be expressed as

$$\begin{aligned} \boldsymbol{\rho}(\vec{P}) &= \frac{1}{2}\left(\mathbf{1} + \vec{P}\cdot\vec{\boldsymbol{\sigma}}\right) \\ &\equiv \frac{1}{2}\left(\mathbf{1} + P_1\boldsymbol{\sigma}_1 + P_2\boldsymbol{\sigma}_2 + P_3\boldsymbol{\sigma}_3\right) \\ &= \frac{1}{2}\begin{pmatrix} 1+P_3 & P_1 - iP_2 \\ P_1 + iP_2 & 1 - P_3 \end{pmatrix}. \end{aligned} \tag{2.62}$$

We can compute $\det\boldsymbol{\rho} = \frac{1}{4}\left(1 - \vec{P}^2\right)$. Therefore, a necessary condition for $\boldsymbol{\rho}$ to have nonnegative eigenvalues is $\det\boldsymbol{\rho} \geq 0$ or $\vec{P}^2 \leq 1$. This condition is also sufficient; since $\mathrm{tr}\boldsymbol{\rho} = 1$, it is not possible for $\boldsymbol{\rho}$ to have two negative eigenvalues. Thus, there is a $1-1$ correspondence between the possible density matrices of a single qubit and the points on the *unit 3-ball* $0 \leq |\vec{P}| \leq 1$. This ball is usually called the *Bloch sphere* (although of course it is really a ball, not a sphere).

The boundary $\left(|\vec{P}| = 1\right)$ of the ball (which really is a sphere) contains the density matrices with vanishing determinant. Since $\mathrm{tr}\boldsymbol{\rho} = 1$, these density matrices must have the eigenvalues 0 and 1. They are one-dimensional projectors, and hence pure states. We have already seen that every pure state of a single qubit is of the form $|\psi(\theta,\varphi)\rangle$ and can be envisioned as a spin pointing in the (θ,φ) direction. Indeed using the property

$$(\hat{n}\cdot\vec{\boldsymbol{\sigma}})^2 = \mathbf{1}, \tag{2.63}$$

where $\hat{n}$ is a unit vector, we can easily verify that the pure-state density matrix

$$\boldsymbol{\rho}(\hat{n}) = \frac{1}{2}\left(\mathbf{1} + \hat{n}\cdot\vec{\boldsymbol{\sigma}}\right) \tag{2.64}$$

satisfies the property

$$(\hat{n}\cdot\vec{\boldsymbol{\sigma}})\,\vec{\boldsymbol{\rho}}(\hat{n}) = \boldsymbol{\rho}(\hat{n})\,(\hat{n}\cdot\vec{\boldsymbol{\sigma}}) = \boldsymbol{\rho}(\hat{n}), \tag{2.65}$$

and, therefore is the projector

$$\boldsymbol{\rho}(\hat{n}) = |\psi(\hat{n})\rangle\langle\psi(\hat{n})|; \tag{2.66}$$

that is, $\hat{n}$ is the direction along which the spin is pointing up. Alternatively, from the expression

$$|\psi(\theta,\phi)\rangle = \begin{pmatrix} e^{-i\varphi/2} & \cos\frac{\theta}{2} \\ e^{i\varphi/2} & \sin\frac{\theta}{2} \end{pmatrix}, \tag{2.67}$$

we may compute directly that

$$\begin{aligned} \rho(\theta,\phi) &= |\psi(\theta,\phi)\rangle\langle\psi(\theta,\phi)| \\ = \begin{pmatrix} \cos^2\frac{\theta}{2} & \cos\frac{\theta}{2}\sin\frac{\theta}{2}e^{-i\varphi} \\ \cos\frac{\theta}{2}\sin\frac{\theta}{2}e^{i\varphi} & \sin^2\frac{\theta}{2} \end{pmatrix} &= \frac{1}{2}\mathbf{1} + \frac{1}{2}\begin{pmatrix} \cos\theta & \sin\theta e^{-i\varphi} \\ \sin\theta e^{i\varphi} & -\cos\theta \end{pmatrix} \\ &= \frac{1}{2}(\mathbf{1} + \hat{n}\cdot\vec{\boldsymbol{\sigma}}) \end{aligned} \tag{2.68}$$

where $\hat{n} = (\sin\theta\cos\varphi, \sin\theta\sin\varphi, \cos\theta)$. One nice property of the Bloch parametrization of the pure states is that while $|\psi(\theta,\varphi)\rangle$ has an arbitrary overall phase that has no physical significance, there is no phase ambiguity in the density matrix $\boldsymbol{\rho}(\theta,\varphi) = |\psi(\theta,\varphi)\rangle\langle\psi(\theta,\varphi)|$; all the parameters in $\boldsymbol{\rho}$ have a physical meaning.

From the property

$$\frac{1}{2}\mathrm{tr}\ \boldsymbol{\sigma}_i\boldsymbol{\sigma}_j = \delta_{ij} \tag{2.69}$$

we see that

$$\langle\hat{n}\cdot\vec{\boldsymbol{\sigma}}\rangle_{\vec{P}} = \mathrm{tr}\left(\hat{n}\cdot\vec{\boldsymbol{\sigma}}\boldsymbol{\rho}(\vec{P})\right) = \hat{n}\cdot\vec{P}\ . \tag{2.70}$$

Thus the vector $\vec{P}$ in Eq. (2.62) parametrizes the *polarization* of the spin. If there are many identically prepared systems at our disposal, we can determine $\vec{P}$ (and hence the complete density matrix $\boldsymbol{\rho}(\vec{P})$) by measuring $\langle\hat{n}\cdot\vec{\boldsymbol{\sigma}}\rangle$ along each of three linearly independent axes.

2.3.3 Gleason's theorem

We arrived at the density matrix $\boldsymbol{\rho}$ and the expression $\mathrm{tr}(\mathbf{M}\boldsymbol{\rho})$ for the expectation value of an observable $\mathbf{M}$ by starting from our axioms of quantum mechanics, and then considering the description of a portion of a larger quantum system. But it is encouraging to know that the density matrix formalism is a very general feature in a much broader framework. This is the content of *Gleason's theorem* (1957).

Gleason's theorem starts from the premise that it is the task of quantum theory to assign consistent probabilities to all possible orthogonal projections in a Hilbert space (in other words, to all possible measurements of observables).

A state of a quantum system, then, is a mapping that take each projection ($\mathbf{E}^2 = \mathbf{E}$ and $\mathbf{E} = \mathbf{E}^\dagger$) to a nonnegative real number less than one:

$$\mathbf{E} \to p(\mathbf{E}); \quad 0 \le p(E) \le 1. \tag{2.71}$$

This mapping must have the properties:

(1) $p(\mathbf{0}) = 0$

(2) $p(\mathbf{1}) = 1$

(3) If $\mathbf{E}_1\mathbf{E}_2 = 0$, then $p(\mathbf{E}_1 + \mathbf{E}_2) = p(\mathbf{E}_1) + p(\mathbf{E}_2)$.

Here (3) is the crucial assumption. It says that (since projections on to mutually orthogonal spaces can be viewed as mutually exclusive alternatives) the probabilities assigned to mutually orthogonal projections must be additive. This assumption is very powerful, because there are so many different ways to choose $\mathbf{E}_1$ and $\mathbf{E}_2$. Roughly speaking, the first two assumptions say that whenever we make a measurement; (1) there is always an outcome, and (2) the probabilities of all possible outcomes sum to 1.

Under these assumptions, Gleason showed that for any such map, there is a hermitian, positive $\boldsymbol{\rho}$ with $\mathrm{tr}\boldsymbol{\rho} = 1$ such that

$$p(\mathbf{E}) = \mathrm{tr}(\boldsymbol{\rho}\mathbf{E}). \tag{2.72}$$

as long as the dimension of the Hilbert space is greater than 2. Thus, the density matrix formalism is really *necessary*, if we are to represent observables as self-adjoint operators in Hilbert space, and we are to consistently assign probabilities to all possible measurement outcomes. Crudely speaking, the requirement of additivity of probabilities for mutually exclusive outcomes is so strong that we are inevitably led to the linear expression eq. (2.72).

The case of a two-dimensional Hilbert space is special because there just are not enough mutually exclusive projections in two dimensions. All nontrivial projections are of the form

$$\mathbf{E}(\hat{n}) = \frac{1}{2}(\mathbf{1} + \hat{n} \cdot \vec{\boldsymbol{\sigma}}), \tag{2.73}$$

and

$$\mathbf{E}(\hat{n})\mathbf{E}(\hat{m}) = 0 \tag{2.74}$$

only for $\hat{m} = -\hat{n}$; therefore, any function $f(\hat{n})$ on the two-sphere such that $f(\hat{n}) + f(-\hat{n}) = 1$ satisfies the premises of Gleason's theorem, and there are many such functions. However, in three-dimensions, there are may more alternative ways to partition unity, so that Gleason's assumptions are far more powerful. The proof of the theorem will not be given here. See Peres, p. 190 ff, for a discussion.

2.3.4 Evolution of the density operator

So far, we have not discussed the time evolution of mixed states. In the case of a bipartite pure state governed by the usual axioms of quantum theory, let us suppose that the Hamiltonian on $\mathcal{H}_A \otimes \mathcal{H}_B$ has the form

$$\mathbf{H}_{AB} = \mathbf{H}_A \otimes \mathbf{1}_B + \mathbf{1}_A \otimes \mathbf{H}_B. \tag{2.75}$$

Under this assumption, there is no coupling between the two subsystems A and B, so that each evolves independently. The time evolution operator for the combined system

$$\mathbf{U}_{AB}(t) = \mathbf{U}_A(t) \otimes \mathbf{U}_B(t), \tag{2.76}$$

decomposes into separate unitary time evolution operators acting on each system.

In the Schrödinger picture of dynamics, then, an initial pure state $|\psi(0)\rangle_{AB}$ of the bipartite system given by eq. (2.57) evolves to

$$|\psi(t)\rangle_{AB} = \sum_{i,\mu} a_{i\mu}|i(t)\rangle_A \otimes |\mu(t)\rangle_B, \tag{2.77}$$

where

$$\begin{aligned} |i(t)\rangle_A &= U_A(t)|i(0)\rangle_A, \\ |\mu(t)\rangle_B &= U_B(t)|\mu(0)\rangle_B, \end{aligned} \tag{2.78}$$

define new orthonormal basis for $\mathcal{H}_A$ and $\mathcal{H}_B$ (since $U_A(t)$ and $U_B(t)$ are unitary). Taking the partial trace as before, we find

$$\begin{aligned} \boldsymbol{\rho}_A(t) &= \sum_{i,j,\mu} a_{i\mu}a^*_{j\nu}|i(t)\rangle_{A\ A}\langle j(t)| \\ &= \mathbf{U}_A(t)\boldsymbol{\rho}_A(0)\mathbf{U}_A(t)^\dagger. \end{aligned} \tag{2.79}$$

Thus $\mathbf{U}_A(t)$, acting by conjugation, determines the time evolution of the density matrix.

In particular, in the basis in which $\boldsymbol{\rho}_A(0)$ is diagonal, we have

$$\boldsymbol{\rho}_A(t) = \sum_a p_a \mathbf{U}_A(t)|\psi_a(0)\rangle_{A\ A}\langle\psi_a(0)|\mathbf{U}_A(t). \tag{2.80}$$

Eq. (2.80) tells us that the evolution of $\boldsymbol{\rho}_A$ is perfectly consistent with the ensemble interpretation. Each state in the ensemble evolves forward in time governed by $\mathbf{U}_A(t)$. If the state $|\psi_a(0)\rangle$ occurs with probability p_a at time 0, then $|\psi_a(t)\rangle$ occurs with probability p_a at the subsequent time t.

On the other hand, it should be clear that eq. (2.80) applies only under the assumption that systems A and B are not *coupled* by the Hamiltonian. Later, we will investigate how the density matrix evolves under more general conditions.

2.4 Schmidt decomposition

A bipartite pure state can be expressed in a standard form (*the Schmidt decomposition*) that is often very useful.

To arrive at this form, note that an arbitrary vector in $\mathcal{H}_A \otimes \mathcal{H}_B$ can be expanded as

$$|\psi\rangle_{AB} = \sum_{i,\mu} a_{i\mu}|i\rangle_A|\mu\rangle_B \equiv \sum_i |i\rangle_A|\tilde{i}\rangle_B. \tag{2.81}$$

Here $\{|i\rangle_A\}$ and $\{|\mu\rangle_B\}$ are orthonormal basis for $\mathcal{H}_A$ and $\mathcal{H}_B$ respectively, but to obtain the second equality in eq. (2.81) we have defined

$$|\tilde{i}\rangle_B \equiv \sum_\mu a_{i\mu}|\mu\rangle_B. \tag{2.82}$$

Note that the $|\tilde{i}\rangle_B$'s need *not* be mutually orthogonal or normalized.

Now let's suppose that the $\{|i\rangle_A\}$ basis is chosen to be the basis in which $\boldsymbol{\rho}_A$ is diagonal,

$$\boldsymbol{\rho}_A = \sum_i p_i|i\rangle_{A\ A}\langle i|. \tag{2.83}$$

We can also compute $\boldsymbol{\rho}_A$ by performing a partial trace,

$$\boldsymbol{\rho}_A = \mathrm{tr}_B(|\psi\rangle_{AB\ AB}\langle\psi|)$$

$$= \mathrm{tr}_B(\sum_{ij} |i\rangle_{A\ A}\langle j| \otimes |\tilde{i}\rangle_{B\ B}\langle\tilde{j}|) = \sum_{ij} {}_B\langle\tilde{j}|\tilde{i}\rangle_B\ (|i\rangle_{A\ A}\langle j|)\ . \tag{2.84}$$

We obtained the last equality in eq. (2.84) by noting that

$$\begin{aligned}\mathrm{tr}_B\left(|\tilde{i}\rangle_{B\ B}\langle\tilde{j}|\right) &= \sum_k\ {}_B\langle k|\tilde{i}\rangle_{B\ B}\langle\tilde{j}|k\rangle_B \\ &= \sum_k\ {}_B\langle\tilde{j}|k\rangle_{B\ B}\langle k|\tilde{i}\rangle_B = {}_B\langle\tilde{j}|\tilde{i}\rangle_B,\end{aligned} \tag{2.85}$$

where $\{|k\rangle_B\}$ is an orthonormal basis for $\mathcal{H}_B$. By comparing eq. (2.83) and eq. (2.84), we see that

$${}_B\langle\tilde{j}|\tilde{i}\rangle_B = p_i\delta_{ij}. \tag{2.86}$$

Hence, it turns out that the $\{|\tilde{i}\rangle_B\}$ *are* orthogonal after all. We obtain orthonormal vectors by rescaling,

$$|i'\rangle_B = p_i^{-1/2}|i\rangle_B \tag{2.87}$$

(we may assume $p_i \neq 0$, because we will need eq. (2.87) only for i appearing in the sum eq. (2.83)), and therefore obtain the expansion

$$|\psi\rangle_{AB} = \sum_i \sqrt{p_i}|i\rangle_A|i'\rangle_B, \tag{2.88}$$

in terms of a *particular* orthonormal basis of $\mathcal{H}_A$ and $\mathcal{H}_B$.

Eq. (2.88) is the Schmidt decomposition of the bipartite pure state $|\psi\rangle_{AB}$. Any bipartite pure state can be expressed in this form, but of course the bases used depend on the pure state that is being expanded. In general, we can't simultaneously expand *both* $|\psi\rangle_{AB}$ and $|\varphi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$ in the form eq. (2.88) using the *same* orthonormal bases for $\mathcal{H}_A$ and $\mathcal{H}_B$.

Using eq. (2.88), we can also evaluate the partial trace over $\mathcal{H}_A$ to obtain

$$\boldsymbol{\rho}_B = \mathrm{tr}_A\left(|\psi\rangle_{AB\ AB}\langle\psi|\right) = \sum_i p_i|i'\rangle_{B\ B}\langle i'|. \tag{2.89}$$

We see that $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ have the *same nonzero eigenvalues*. Of course, there is no need for $\mathcal{H}_A$ and $\mathcal{H}_B$ to have the same dimension, so the number of *zero* eigenvalues of $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ can differ.

If $\boldsymbol{\rho}_A$ (and hence $\boldsymbol{\rho}_B$) have no degenerate eigenvalues other than zero, then the Schmidt decomposition of $|\psi\rangle_{AB}$ is essentially uniquely determined by $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$. We can diagonalize $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ to find the $|i\rangle_A$'s and $|i'\rangle_B$'s, and then we pair up the eigenstates of $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ with the same eigenvalue to obtain eq. (2.88). We have chosen the phases of our basis states so that no phases appear in the coefficients in the sum; the only remaining freedom is to redefine $|i\rangle_A$ and $|i'\rangle_B$ by multiplying by opposite phases (which of course leaves the expression eq. (2.88) unchanged).

But if $\boldsymbol{\rho}_A$ has degenerate nonzero eigenvalues, then we need more information than that provided by $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ to determine the Schmidt decomposition; we need to know which $|i'\rangle_B$ gets paired with each $|i\rangle_A$. For example, if both $\mathcal{H}_A$ and $\mathcal{H}_B$ are N-dimensional and $\mathbf{U}_{ij}$ is any $N \times N$ unitary matrix, then

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{N}} \sum_{i,j=1}^{N} |i\rangle_A \mathbf{U}_{ij} |j'\rangle_B, \tag{2.90}$$

will yield $\boldsymbol{\rho}_A = \boldsymbol{\rho}_B = \frac{1}{N}\mathbf{1}$ when we take partial traces. Furthermore, we are free to apply simultaneous unitary transformations in $\mathcal{H}_A$ and $\mathcal{H}_B$,

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{N}} \sum_{i} |i\rangle_A |i'\rangle_B = \frac{1}{\sqrt{N}} \sum_{ijk} \mathbf{U}^*_{ij} |j\rangle_A \mathbf{U}_{ik} |k'\rangle_B; \tag{2.91}$$

this preserves the state $|\psi\rangle_{AB}$, but illustrates that there is an ambiguity in the basis used when we express $|\psi\rangle_{AB}$ in the Schmidt form.

2.4.1 Entanglement

With any bipartite pure state $|\psi\rangle_{AB}$ we may associate a positive integer, the *Schmidt number*, which is the number of nonzero eigenvalues in $\boldsymbol{\rho}_A$ (or $\boldsymbol{\rho}_B$) and hence the number of terms in the Schmidt decomposition of $|\psi\rangle_{AB}$. In terms of this quantity, we can define what it means for a bipartite pure state to be *entangled*: $|\psi\rangle_{AB}$ is entangled (or nonseparable) if its Schmidt number is greater than one; otherwise, it is *separable* (or unentangled). Thus, a separable bipartite pure state is a direct product of pure states in $\mathcal{H}_A$ and $\mathcal{H}_B$,

$$|\psi\rangle_{AB} = |\varphi\rangle_A \otimes |\chi\rangle_B; \tag{2.92}$$

then the reduced density matrices $\boldsymbol{\rho}_A = |\varphi\rangle_A\, {}_A\langle\varphi|$ and $\boldsymbol{\rho}_B = |\chi\rangle_B\, {}_B\langle\chi|$ are pure. Any state that cannot be expressed as such a direct product is entangled; then $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ are mixed states.

One of our main goals this term will be to understand better the significance of entanglement. It is not strictly correct to say that subsystems A and B are *uncorrelated* if $|\psi\rangle_{AB}$ is separable; after all, the two spins in the separable state

$$|\uparrow\rangle_A |\uparrow\rangle_B, \tag{2.93}$$

are surely correlated – they are both pointing in the same direction. But the correlations between A and B in an entangled state have a different character than those in a separable state. Perhaps the critical difference is that *entanglement cannot be created locally.* The only way to entangle A and B is for the two subsystems to directly interact with one another.

We can prepare the state eq. (2.93) without allowing spins A and B to ever come into contact with one another. We need only send a (classical!) message to two preparers (Alice and Bob) telling both of them to prepare a spin pointing along the z-axis. But the only way to turn the state eq. (2.93) into an entangled state like

$$\frac{1}{\sqrt{2}} \left(|\uparrow\rangle_A |\uparrow\rangle_B + |\downarrow\rangle_A |\downarrow\rangle_B \right), \tag{2.94}$$

is to apply a *collective* unitary transformation to the state. Local unitary transformations of the form $U_A \otimes U_B$, and local measurements performed by Alice or Bob, *cannot increase the Schmidt number* of the two-qubit state, no matter how much Alice and Bob discuss what they do. To entangle two qubits, we *must* bring them together and allow them to interact.

As we will discuss later, it is also possible to make the distinction between entangled and separable bipartite *mixed* states. We will also discuss various ways in which local operations can modify the form of entanglement, and some ways that entanglement can be put to use.

2.5 Ambiguity of the ensemble interpretation

2.5.1 Convexity

Recall that an operator $\boldsymbol{\rho}$ acting on a Hilbert space $\mathcal{H}$ may be interpreted as a density operator if it has the three properties:

(1) $\boldsymbol{\rho}$ is self-adjoint.

(2) $\boldsymbol{\rho}$ is nonnegative.

(3) $\mathrm{tr}(\boldsymbol{\rho}) = 1$.

It follows immediately that, given two density matrices $\boldsymbol{\rho}_1$, and $\boldsymbol{\rho}_2$, we can always construct another density matrix as a convex linear combination of the two:

$$\boldsymbol{\rho}(\lambda) = \lambda\boldsymbol{\rho}_1 + (1-\lambda)\boldsymbol{\rho}_2 \tag{2.95}$$

is a density matrix for any real λ satisfying $0 \le \lambda \le 1$. We easily see that $\boldsymbol{\rho}(\lambda)$ satisfies (1) and (3) if $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$ do. To check (2), we evaluate

$$\langle\psi|\boldsymbol{\rho}(\lambda)|\psi\rangle = \lambda\langle\psi|\boldsymbol{\rho}_1|\psi\rangle + (1-\lambda)\langle\psi|\boldsymbol{\rho}_2|\psi\rangle \ge 0; \tag{2.96}$$

$\langle\boldsymbol{\rho}(\lambda)\rangle$ is guaranteed to be nonnegative because $\langle\boldsymbol{\rho}_1\rangle$ and $\langle\boldsymbol{\rho}_2\rangle$ are. We have, therefore, shown that in a Hilbert space $\mathcal{H}$ of dimension N, the density operators are a *convex subset* of the real vector space of $N \times N$ hermitian matrices. (A subset of a vector space is said to be convex if the set contains the straight line segment connecting any two points in the set.)

Most density operators can be expressed as a sum of other density operators in many different ways. But the pure states are special in this regard – it is *not* possible to express a pure state as a convex sum of two other states. Consider a pure state $\boldsymbol{\rho} = |\psi\rangle\langle\psi|$, and let $|\psi_\perp\rangle$ denote a vector orthogonal to $|\psi\rangle$, $\langle\psi_\perp|\psi\rangle = 0$. Suppose that $\boldsymbol{\rho}$ can be expanded as in eq. (2.95); then

$$\begin{aligned}\langle\psi_\perp|\boldsymbol{\rho}|\psi_\perp\rangle = 0 = \lambda&\langle\psi_\perp|\boldsymbol{\rho}_1|\psi_\perp\rangle \\ &+ (1-\lambda)\langle\psi_\perp|\boldsymbol{\rho}_2|\psi_\perp\rangle.\end{aligned} \tag{2.97}$$

Since the right hand side is a sum of two nonnegative terms, and the sum vanishes, both terms must vanish. If λ is not 0 or 1, we conclude that $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$ are orthogonal to $|\psi_\perp\rangle$. But since $|\psi_\perp\rangle$ can be *any* vector orthogonal to $|\psi\rangle$, we conclude that $\boldsymbol{\rho}_1 = \boldsymbol{\rho}_2 = \boldsymbol{\rho}$.

The vectors in a convex set that cannot be expressed as a linear combination of other vectors in the set are called the *extremal points* of the set. We have just shown that the pure states are extremal points of the set of density matrices. Furthermore, *only* the pure states are extremal, because any mixed state can be written $\boldsymbol{\rho} = \sum_i p_i|i\rangle\langle i|$ in the basis in which it is diagonal, and so is a convex sum of pure states.

We have already encountered this structure in our discussion of the special case of the Bloch sphere. We saw that the density operators are a (unit) ball in the three-dimensional set of 2×2 hermitian matrices with unit trace. The ball is convex, and its extremal points are the points on the boundary. Similarly, the $N \times N$ density operators are a convex subset of the (N^2-1)-dimensional set of $N \times N$ hermitian matrices with unit trace, and the extremal points of the set are the pure states.

However, the 2×2 case is atypical in one respect: for $N > 2$, the points on the boundary of the set of density matrices are not necessarily pure states. The boundary of the set consists of all density matrices with at least one vanishing eigenvalue (since there are nearby matrices with negative eigenvalues). Such a density matrix need not be pure, for $N > 2$, since the number of nonvanishing eigenvalues can exceed one.

2.5.2 Ensemble preparation

The convexity of the set of density matrices has a simple and enlightening physical interpretation. Suppose that a preparer agrees to prepare one of two possible states; with probability λ, the state $\boldsymbol{\rho}_1$ is prepared, and with probability $1-\lambda$, the state $\boldsymbol{\rho}_2$ is prepared. (A random number generator might be employed to guide this choice.) To evaluate the expectation value of any observable $\mathbf{M}$, we average over *both* the choices of preparation *and* the outcome of the quantum measurement:

$$\begin{aligned}\langle\mathbf{M}\rangle &= \lambda\langle\mathbf{M}\rangle_1 + (1-\lambda)\langle\mathbf{M}\rangle_2 \\ &= \lambda\mathrm{tr}(\mathbf{M}\boldsymbol{\rho}_1) + (1-\lambda)\mathrm{tr}(\mathbf{M}\boldsymbol{\rho}_2) \\ &= \mathrm{tr}\left(\mathbf{M}\boldsymbol{\rho}(\lambda)\right).\end{aligned} \tag{2.98}$$

All expectation values are thus indistinguishable from what we would obtain if the state $\boldsymbol{\rho}(\lambda)$ had been prepared instead. Thus, we have an operational procedure, given methods for preparing the states $\boldsymbol{\rho}_1$ and $\boldsymbol{\rho}_2$, for preparing any convex combination.

Indeed, for any mixed state $\boldsymbol{\rho}$, there are an infinite variety of ways to express $\boldsymbol{\rho}$ as a convex combination of other states, and hence an infinite variety of procedures we could employ to prepare $\boldsymbol{\rho}$, all of which have exactly the same consequences for any conceivable observation of the system. But a pure state is different; it can be prepared in only one way. (This is what is "pure" about a pure state.) Every pure state is an eigenstate of some observable, e.g., for the state $\boldsymbol{\rho} = |\psi\rangle\langle\psi|$, measurement of the projection $\mathbf{E} = |\psi\rangle\langle\psi|$ is guaranteed to have the outcome 1. (For example, recall that every pure state of a single qubit is "spin-up" along some axis.) Since $\boldsymbol{\rho}$ is the only state for which the outcome of measuring $\mathbf{E}$ is 1 with 100% probability, there is no way to reproduce this observable property by choosing one of several possible preparations. Thus, the preparation of a pure state is unambiguous (we can determine a unique preparation if we have many copies of the state to experiment with), but the preparation of a mixed state is always ambiguous.

How ambiguous is it? Since any $\boldsymbol{\rho}$ can be expressed as a sum of pure states, let's confine our attention to the question: in how many ways can a density operator be expressed as a convex sum of pure states? Mathematically, this is the question: in how many ways can $\boldsymbol{\rho}$ be written as a sum of *extremal* states?

As a first example, consider the "maximally mixed" state of a single qubit:

$$\boldsymbol{\rho} = \frac{1}{2}\mathbf{1}. \tag{2.99}$$

This can indeed be prepared as an ensemble of pure states in an infinite variety of ways. For example,

$$\boldsymbol{\rho} = \frac{1}{2}|\uparrow_z\rangle\langle\uparrow_z| + \frac{1}{2}|\downarrow_z\rangle\langle\downarrow_z|, \tag{2.100}$$

so we obtain $\boldsymbol{\rho}$ if we prepare either $|\uparrow_z\rangle$ or $|\downarrow_z\rangle$, each occurring with probability $\frac{1}{2}$. But we also have

$$\boldsymbol{\rho} = \frac{1}{2}|\uparrow_x\rangle\langle\uparrow_x| + \frac{1}{2}|\downarrow_x\rangle\langle\downarrow_x|, \tag{2.101}$$

so we obtain $\boldsymbol{\rho}$ if we prepare either $|\uparrow_x\rangle$ or $|\downarrow_x\rangle$, each occurring with probability $\frac{1}{2}$. Now the preparation procedures are undeniably *different*. Yet there is no possible way to tell the difference by making observations of the spin.

More generally, the point at the center of the Bloch ball is the sum of any two antipodal points on the sphere – preparing either $|\uparrow_{\hat{n}}\rangle$ or $|\downarrow_{\hat{n}}\rangle$, each occurring with probability $\frac{1}{2}$ will generate $\boldsymbol{\rho} = \frac{1}{2}\mathbf{1}$.

Only in the case where $\boldsymbol{\rho}$ has two (or more) degenerate eigenvalues will there be distinct ways of generating $\boldsymbol{\rho}$ from an ensemble of *mutually orthogonal* pure states, but there is no good reason to confine our attention to ensembles of mutually orthogonal pure states. We may consider a point in the interior of the Bloch ball

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2}(\mathbf{1} + \vec{P}\cdot\vec{\sigma}), \tag{2.102}$$

with $0 < |\vec{P}| < 1$, and it too can be expressed as

$$\boldsymbol{\rho}(\vec{P}) = \lambda\boldsymbol{\rho}(\hat{n}_1) + (1-\lambda)\boldsymbol{\rho}(\hat{n}_2), \tag{2.103}$$

if $\vec{P} = \lambda\hat{n}_1 + (1-\lambda)\hat{n}_2$ (or in other words, if $\vec{P}$ lies somewhere on the line segment connecting the points $\hat{n}_1$ and $\hat{n}_2$ on the sphere). Evidently, for any $\vec{P}$, there is a solution associated with any chord of the sphere that passes through the point $\vec{P}$; all such chords comprise a two-parameter family.

This highly ambiguous nature of the preparation of a mixed quantum state is one of the characteristic features of quantum information that contrasts sharply with classical probability distributions. Consider, for example, the case of a probability distribution for a single classical bit. The two extremal distributions are those in which either 0 or 1 occurs with 100% probability. *Any* probability distribution for the bit is a convex sum of these two extremal points. Similarly, if there are N possible states, there are N extremal distributions, and any probability distribution has a *unique* decomposition into extremal ones (the convex set of probability distributions is a *simplex*). If 0 occurs with 21% probability, 1 with 33% probability, and 2 with 46% probability, there is a unique preparation procedure that yields this probability distribution!

2.5.3 Faster than light?

Let's now return to our earlier viewpoint – that a mixed state of system A arises because A is *entangled* with system B – to further consider the implications of the ambiguous preparation of mixed states. If qubit A has density matrix

$$\boldsymbol{\rho}_A = \frac{1}{2}|\uparrow_z\rangle_A\,{}_A\langle\uparrow_z| + \frac{1}{2}|\downarrow_z\rangle_A\,{}_A\langle\downarrow_z|, \tag{2.104}$$

this density matrix could arise from an entangled bipartite pure state $|\psi\rangle_{AB}$ with the Schmidt decomposition

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{2}}\left(|\uparrow_z\rangle_A|\uparrow_z\rangle_B + |\downarrow_z\rangle_A|\downarrow_z\rangle_B\right). \tag{2.105}$$

Therefore, the ensemble interpretation of $\boldsymbol{\rho}_A$ in which either $|\uparrow_z\rangle_A$ or $|\downarrow_z\rangle_A$ is prepared (each with probability $p=\frac{1}{2}$) can be realized by performing a measurement of qubit B. We measure qubit B in the $\{|\uparrow_z\rangle_B, |\downarrow_z\rangle_B\}$ basis; if the result $|\uparrow_z\rangle_B$ is obtained, we have prepared $|\uparrow_z\rangle_A$, and if the result $|\downarrow_z\rangle_B$ is obtained, we have prepared $|\downarrow_z\rangle_A$.

But as we have already noted, in this case, because $\boldsymbol{\rho}_A$ has degenerate eigenvalues, the Schmidt basis is not unique. We can apply simultaneous unitary transformations to qubits A and B (actually, if we apply U to A we must apply U^* to B) without modifying the bipartite pure state $|\psi\rangle_{AB}$. Therefore, for *any* unit 3-vector $\hat{n}$, $|\psi\rangle_{AB}$ has a Schmidt decomposition of the form

$$|\psi\rangle_{AB} = \frac{1}{\sqrt{2}}\left(|\uparrow_{\hat{n}}\rangle_A |\uparrow_{\hat{n}'}\rangle_B + |\downarrow_{\hat{n}}\rangle_A |\downarrow_{\hat{n}'}\rangle_B\right). \tag{2.106}$$

We see that by measuring qubit B in a suitable basis, we can realize *any* interpretation of $\boldsymbol{\rho}_A$ as an ensemble of two pure states.

Bright students, upon learning of this property, are sometimes inspired to suggest a mechanism for faster-than-light communication. Many copies of $|\psi\rangle_{AB}$ are prepared. Alice takes all of the A qubits to the Andromeda galaxy and Bob keeps all of the B qubits on earth. When Bob wants to send a one-bit message to Alice, he chooses to measure either $\boldsymbol{\sigma}_1$ or $\boldsymbol{\sigma}_3$ for all his spins, thus preparing Alice's spins in either the $\{|\uparrow_z\rangle_A, |\downarrow_z\rangle_A\}$ or $\{|\uparrow_x\rangle_A, |\downarrow_x\rangle_A\}$ ensembles.[1] To read the message, Alice immediately measures her spins to see which ensemble has been prepared.

But *exceptionally* bright students (or students who heard the previous lecture) can see the flaw in this scheme. Though the two preparation methods are surely different, both ensembles are described by precisely the same density matrix $\boldsymbol{\rho}_A$. Thus, there is no conceivable measurement Alice can make that will distinguish the two ensembles, and no way for Alice to tell what action Bob performed. The "message" is unreadable.

Why, then, do we confidently state that "the two preparation methods are surely different?" To qualm any doubts about that, imagine that Bob either (1) measures all of his spins along the $\hat{z}$-axis, or (2) measures all of his spins along the $\hat{x}$-axis, and then calls Alice on the intergalactic telephone. He does *not* tell Alice whether he did (1) or (2), but he does tell her the results of all his measurements: "the first spin was up, the second was down," etc. Now

[1] U is real in this case, so $U = U^*$ and $\hat{n} = \hat{n}'$.

Alice performs either (1) or (2) on *her* spins. If both Alice and Bob measured along the same axis, Alice will find that every single one of her measurement outcomes agrees with what Bob found. But if Alice and Bob measured along different (orthogonal) axes, then Alice will find *no correlation* between her results and Bob's. About half of her measurements agree with Bob's and about half disagree. If Bob promises to do either (1) or (2), and assuming no preparation or measurement errors, then Alice will know that Bob's action was different than hers (even though Bob never told her this information) as soon as one of her measurements disagrees with what Bob found. If all their measurements agree, then if many spins are measured, Alice will have very high statistical confidence that she and Bob measured along the same axis. (Even with occasional measurement errors, the statistical test will still be highly reliable if the error rate is low enough.) So Alice does have a way to distinguish Bob's two preparation methods, but in this case there is certainly no faster-than-light communication, because Alice had to receive Bob's phone call before she could perform her test.

2.5.4 Quantum erasure

We had said that the density matrix $\boldsymbol{\rho}_A = \frac{1}{2}\mathbf{1}$ describes a spin in an *incoherent* superposition of the pure states $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. This was to be distinguished from *coherent* superpositions of these states, such as

$$|\uparrow_x, \downarrow_x\rangle = \frac{1}{2}\left(|\uparrow_z\rangle \pm |\downarrow_z\rangle\right) ; \tag{2.107}$$

in the case of a coherent superposition, the *relative phase* of the two states has observable consequences (distinguishes $|\uparrow_x\rangle$ from $|\downarrow_x\rangle$). In the case of an incoherent superposition, the relative phase is completely unobservable. The superposition becomes incoherent if spin A becomes entangled with another spin B, and spin B is inaccessible.

Heuristically, the states $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ can *interfere* (the relative phase of these states can be observed) only if we have no information about whether the spin state is $|\uparrow_z\rangle_A$ or $|\downarrow_z\rangle_A$. More than that, interference can occur only if there is *in principle no possible way* to find out whether the spin is up or down along the z-axis. Entangling spin A with spin B destroys interference, (causes spin A to *decohere*) because it is possible in principle for us to determine if spin A is up or down along $\hat{z}$ by performing a suitable measurement of spin B.

But we have now seen that the statement that entanglement causes decoherence requires a qualification. Suppose that Bob measures spin B along the $\hat{x}$-axis, obtaining either the result $|\uparrow_x\rangle_B$ or $|\downarrow_x\rangle_B$, and that he sends his measurement result to Alice. *Now* Alice's spin is a pure state (either $|\uparrow_x\rangle_A$ or $|\downarrow_x\rangle_A$) and in fact a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. We have managed to recover the purity of Alice's spin before the jaws of decoherence could close!

Suppose that Bob allows his spin to pass through a Stern–Gerlach apparatus oriented along the $\hat{z}$-axis. Well, of course, Alice's spin can't behave like a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$; all Bob has to do is look to see which way his spin moved, and he will know whether Alice's spin is up or down along $\hat{z}$. But suppose that Bob does not look. Instead, he carefully refocuses the two beams without maintaining any record of whether his spin moved up or down, and *then* allows the spin to pass through a second Stern–Gerlach apparatus oriented along the $\hat{x}$-axis. *This* time he looks, and communicates the result of his $\boldsymbol{\sigma}_1$ measurement to Alice. Now the coherence of Alice's spin has been restored!

This situation has been called a *quantum eraser*. Entangling the two spins creates a "measurement situation" in which the coherence of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ is lost because we can find out if spin A is up or down along $\hat{z}$ by observing spin B. But when we measure spin B along $\hat{x}$, this information is "erased." Whether the result is $|\uparrow_x\rangle_B$ or $|\downarrow_x\rangle_B$ does not tell us anything about whether spin A is up or down along $\hat{z}$, because Bob has been careful not to retain the "which way" information that he might have acquired by looking at the first Stern–Gerlach apparatus.[2] Therefore, it is possible again for spin A to behave like a coherent superposition of $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$ (and it does, *after* Alice hears about Bob's result).

We can best understand the quantum eraser from the ensemble viewpoint. Alice has many spins selected from an ensemble described by $\boldsymbol{\rho}_A = \frac{1}{2}\mathbf{1}$, and there is no way for her to observe interference between $|\uparrow_z\rangle_A$ and $|\downarrow_z\rangle_A$. When Bob makes his measurement along $\hat{x}$, a particular preparation of the ensemble is realized. However, this has no effect that Alice can perceive – her spin is *still* described by $\boldsymbol{\rho}_A = \frac{1}{2}\mathbf{1}$ as before. But, when Alice receives Bob's phone call, she can select a *subensemble* of her spins that are all in the pure state $|\uparrow_x\rangle_A$. The information that Bob sends allows Alice to distill

[2]One often says that the "welcher weg" information has been erased, because it sounds more sophisticated in German.

purity from a maximally mixed state.

Another wrinkle on the quantum eraser is sometimes called *delayed choice*. This just means that the situation we have described is really completely symmetric between Alice and Bob, so it can't make any difference who measures first. (Indeed, if Alice's and Bob's measurements are spacelike separated events, there is no invariant meaning to which came first; it depends on the frame of reference of the observer.) Alice could measure all of her spins today (say along $\hat{x}$) before Bob has made his mind up how he will measure his spins. Next week, Bob can decide to "prepare" Alice's spins in the states $|\uparrow_{\hat{n}}\rangle_A$ and $|\downarrow_{\hat{n}}\rangle_A$ (that is the "delayed choice"). He then tells Alice which were the $|\uparrow_{\hat{n}}\rangle_A$ spins, and she can check her measurement record to verify that

$$\langle\sigma_1\rangle_{\hat{n}} = \hat{n}\cdot\hat{x} \ . \tag{2.108}$$

The results are the same, irrespective of whether Bob "prepares" the spins before or after Alice measures them.

We have claimed that the density matrix $\boldsymbol{\rho}_A$ provides a complete physical description of the state of subsystem A, because it characterizes all possible measurements that can be performed on A. One sometimes hears the objection[3] that the quantum eraser phenomenon demonstrates otherwise. Since the information received from Bob enables Alice to recover a pure state from the mixture, how can we hold that everything Alice can know about A is encoded in $\boldsymbol{\rho}_A$?

I don't think this is the right conclusion. Rather, I would say that quantum erasure provides yet another opportunity to recite our mantra: "Information is physical." The state $\boldsymbol{\rho}_A$ of system A is not the same thing as $\boldsymbol{\rho}_A$ accompanied by the information that Alice has received from Bob. This information (which attaches labels to the subensembles) changes the physical description. One way to say this mathematically is that we should include Alice's "state of knowledge" in our description. An ensemble of spins for which Alice has no information about whether each spin is up or down is a *different* physical state than an ensemble in which Alice knows which spins are up and which are down.[4]

[3]For example, from Roger Penrose in *Shadows of the Mind.*

[4]This "state of knowledge" need not really be the state of a human mind; any (inanimate) record that labels the subensemble will suffice.

2.5.5 The GHJW theorem

So far, we have considered the quantum eraser only in the context of a single qubit, described by an ensemble of equally probable mutually orthogonal states, ($i.e.$, $\boldsymbol{\rho}_A = \frac{1}{2}\mathbf{1}$). The discussion can be considerably generalized.

We have already seen that a mixed state of any quantum system can be realized as an ensemble of pure states in an infinite number of different ways. For a density matrix $\boldsymbol{\rho}_A$, consider one such realization:

$$\boldsymbol{\rho}_A = \sum_i p_i |\varphi_i\rangle_A \, {}_A\langle\varphi_i|, \quad \sum p_i = 1. \tag{2.109}$$

Here the states $\{|\varphi_i\rangle_A\}$ are all normalized vectors, but we do *not* assume that they are mutually orthogonal. Nevertheless, $\boldsymbol{\rho}_A$ can be realized as an ensemble, in which each pure state $|\varphi_i\rangle_A \, {}_A\langle\varphi_i|$ occurs with probability p_i.

Of course, for any such $\boldsymbol{\rho}_A$, we can construct a "purification" of $\boldsymbol{\rho}_A$, a bipartite pure state $|\Phi_1\rangle_{AB}$ that yields $\boldsymbol{\rho}_A$ when we perform a partial trace over $\mathcal{H}_B$. One such purification is of the form

$$|\Phi_1\rangle_{AB} = \sum_i \sqrt{p_i}|\varphi_i\rangle_A|\alpha_i\rangle_B, \tag{2.110}$$

where the vectors $|\alpha_i\rangle_B \in \mathcal{H}_B$ *are* mutually orthogonal and normalized,

$${}_B\langle\alpha_i|\alpha_j\rangle_B = \delta_{ij}. \tag{2.111}$$

Clearly, then,

$$\mathrm{tr}_B\left(|\Phi_1\rangle_{AB} \, {}_{AB}\langle\Phi_1|\right) = \boldsymbol{\rho}_A. \tag{2.112}$$

Furthermore, we can imagine performing an orthogonal measurement in system B that projects onto the $|\alpha_i\rangle_B$ basis.[5] The outcome $|\alpha_i\rangle_B$ will occur with probability p_i, and will prepare the pure state $|\varphi_i\rangle_A \, {}_A\langle\varphi_i|$ of system A. Thus, given the purification $|\Phi\rangle_{AB}$ of $\boldsymbol{\rho}_A$, there is a measurement we can perform in system B that realizes the $|\varphi_i\rangle_A$ ensemble interpretation of $\boldsymbol{\rho}_A$. When the measurement outcome in B is known, we have successfully extracted one of the pure states $|\varphi_i\rangle_A$ from the mixture $\boldsymbol{\rho}_A$.

What we have just described is a generalization of preparing $|\uparrow_z\rangle_A$ by measuring spin B along $\hat{z}$ (in our discussion of two entangled qubits). But

[5]The $|\alpha_i\rangle_B$'s might not span $\mathcal{H}_B$, but in the state $|\Phi\rangle_{AB}$, measurement outcomes orthogonal to all the $|\alpha_i\rangle_B$'s never occur.

to generalize the notion of a quantum eraser, we wish to see that in the state $|\Phi_1\rangle_{AB}$, we can realize a *different* ensemble interpretation of $\boldsymbol{\rho}_A$ by performing a different measurement of B. So let

$$\boldsymbol{\rho}_A = \sum_\mu q_\mu |\psi_\mu\rangle_A \, {}_A\langle\psi_\mu|, \tag{2.113}$$

be another realization of the same density matrix $\boldsymbol{\rho}_A$ as an ensemble of pure states. For this ensemble as well, there is a corresponding purification

$$|\Phi_2\rangle_{AB} = \sum_\mu \sqrt{q_\mu}|\psi_\mu\rangle_A \otimes |\beta_\mu\rangle_B, \tag{2.114}$$

where again the $\{|\beta_\mu\rangle_B$'s$\}$ are orthonormal vectors in $\mathcal{H}_B$. So in the state $|\Phi_2\rangle_{AB}$, we can realize the ensemble by performing a measurement in $\mathcal{H}_B$ that projects onto the $\{|\beta_\mu\rangle_B\}$ basis.

Now, how are $|\Phi_1\rangle_{AB}$ and $|\Phi_2\rangle_{AB}$ related? In fact, we can easily show that

$$|\Phi_1\rangle_{AB} = (\mathbf{1}_A \otimes \mathbf{U}_B)\,|\Phi_2\rangle_{AB}; \tag{2.115}$$

the two states differ by a unitary change of basis acting in $\mathcal{H}_B$ alone, or

$$|\Phi_1\rangle_{AB} = \sum_\mu \sqrt{q_\mu}|\psi_\mu\rangle_A|\gamma_\mu\rangle_B, \tag{2.116}$$

where

$$|\gamma_\mu\rangle_B = \mathbf{U}_B|\beta_\mu\rangle_B, \tag{2.117}$$

is yet another orthonormal basis for $\mathcal{H}_B$. We see, then, that there is a *single* purification $|\Phi_1\rangle_{AB}$ of $\boldsymbol{\rho}_A$, such that we can realize either the $\{|\varphi_i\rangle_A\}$ ensemble or $\{|\psi_\mu\rangle_A\}$ ensemble by choosing to measure the appropriate observable in system B!

Similarly, we may consider many ensembles that all realize $\boldsymbol{\rho}_A$, where the maximum number of pure states appearing in any of the ensembles is n. Then we may choose a Hilbert space $\mathcal{H}_B$ of dimension n, and a pure state $|\Phi\rangle_{AB} \in \mathcal{H}_A \otimes \mathcal{H}_B$, such that any one of the ensembles can be realized by measuring a suitable observable of B. This is the *GHJW*[6] *theorem*. It expresses the quantum eraser phenomenon in its most general form.

[6]For Gisin and Hughston, Jozsa, and Wootters.

In fact, the GHJW theorem is an almost trivial corollary to the Schmidt decomposition. Both $|\Phi_1\rangle_{AB}$ and $|\Phi_2\rangle_{AB}$ have a Schmidt decomposition, and because both yield the same $\boldsymbol{\rho}_A$ when we take the partial trace over B, these decompositions must have the form

$$|\Phi_1\rangle_{AB} = \sum_k \sqrt{\lambda_k}|k\rangle_A|k'_1\rangle_B,$$
$$|\Phi_2\rangle_{AB} = \sum_k \sqrt{\lambda_k}|k\rangle_A|k'_2\rangle_B, \quad (2.118)$$

where the λ_k's are the eigenvalues of $\boldsymbol{\rho}_A$ and the $|k\rangle_A$'s are the corresponding eigenvectors. But since $\{|k'_1\rangle_B\}$ and $\{|k'_2\rangle_B\}$ are both orthonormal bases for $\mathcal{H}_B$, there is a unitary $\mathbf{U}_B$ such that

$$|k'_1\rangle_B = \mathbf{U}_B|k'_2\rangle_B, \quad (2.119)$$

from which eq. (2.115) immediately follows.

In the ensemble of pure states described by Eq. (2.109), we would say that the pure states $|\varphi_i\rangle_A$ are superposed *incoherently* — an observer in system A cannot detect the relative phases of these states. Heuristically, the reason that these states cannot interfere is that it is possible in principle to find out which representative of the ensemble is actually realized by performing a measurement in system B, a projection onto the orthonormal basis $\{|\alpha_i\rangle_B\}$. However, by projecting onto the $\{|\gamma_\mu\rangle_B\}$ basis instead, and relaying the information about the measurement outcome to system A, we can extract one of the pure states $|\psi_\mu\rangle_A$ from the ensemble, even though this state may be a coherent superposition of the $|\varphi_i\rangle_A$'s. In effect, measuring B in the $\{|\gamma_\mu\rangle_B\}$ basis "erases" the "welcher weg" information (whether the state of A is $|\varphi_i\rangle_A$ or $|\varphi_j\rangle_A$). In this sense, the GHJW theorem characterizes the general quantum eraser. The moral, once again, is that *information is physical* — the information acquired by measuring system B, when relayed to A, changes the physical description of a state of A.

2.6 Summary

Axioms. The arena of quantum mechanics is a Hilbert space $\mathcal{H}$. The fundamental assumptions are:

(1) A *state* is a *ray* in $\mathcal{H}$.

(2) An *observable* is a *self-adjoint operator* on $\mathcal{H}$.

(3) A *measurement* is an orthogonal *projection*.

(4) *Time evolution* is *unitary*.

Density operator. But if we confine our attention to only a portion of a larger quantum system, assumptions (1)-(4) need not be satisfied. In particular, a quantum state is described not by a ray, but by a density operator $\boldsymbol{\rho}$, a nonnegative operator with unit trace. The density operator is *pure* (and the state can be described by a ray) if $\boldsymbol{\rho}^2 = \boldsymbol{\rho}$; otherwise, the state is *mixed*. An observable $\mathbf{M}$ has expectation value $\mathrm{tr}(\mathbf{M}\boldsymbol{\rho})$ in this state.

Qubit. A quantum system with a two-dimensional Hilbert space is called a *qubit*. The general density matrix of a qubit is

$$\boldsymbol{\rho}(\vec{P}) = \frac{1}{2}(\mathbf{1} + \vec{P}\cdot\vec{\boldsymbol{\sigma}}) \quad (2.120)$$

where $\vec{P}$ is a three-component vector of length $|\vec{P}| \leq 1$. Pure states have $|\vec{P}| = 1$.

Schmidt decomposition. For any quantum system divided into two parts A and B (a *bipartite* system), the Hilbert space is a tensor product $\mathcal{H}_A \oplus \mathcal{H}_B$. For any pure state $|\psi\rangle_{AB}$ of a bipartite system, there are orthonormal bases $\{|i\rangle_A\}$ for $\mathcal{H}_A$ and $\{|i'\rangle_B\}$ for $\mathcal{H}_B$ such that

$$|\psi\rangle_{AB} = \sum_i \sqrt{p_i}|i\rangle_A|i'\rangle_B; \quad (2.121)$$

Eq. (2.121) is called the *Schmidt decomposition* of $|\psi\rangle_{AB}$. In a bipartite pure state, subsystems A and B separately are described by density operators $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$; it follows from eq. (2.121) that $\boldsymbol{\rho}_A$ and $\boldsymbol{\rho}_B$ have the same nonvanishing eigenvalues (the p_i's). The number of nonvanishing eigenvalues is called the *Schmidt number* of $|\psi\rangle_{AB}$. A bipartite pure state is said to be *entangled* if its Schmidt number is greater than one.

Ensembles. The density operators on a Hilbert space form a convex set, and the pure states are the *extremal points* of the set. A mixed state of a system A can be prepared as an *ensemble* of pure states in many different ways, all of which are experimentally indistinguishable if we observe system A alone. Given any mixed state $\boldsymbol{\rho}_A$ of system A, any preparation of $\boldsymbol{\rho}_A$ as an ensemble of pure states can be realized in principle by performing a measurement in another system B with which A is entangled. In fact given many such preparations of $\boldsymbol{\rho}_A$, there is a single entangled state of A and B such that any one of these preparations can be realized by measuring a suitable observable in B (the *GHJW theorem*). By measuring in system B and reporting the measurement outcome to system A, we can extract from the mixture a pure state chosen from one of the ensembles.

2.7 Exercises

2.1 Fidelity of a random guess

A single qubit (spin-$\frac{1}{2}$ object) is in an unknown *pure* state $|\psi\rangle$, selected at random from an ensemble uniformly distributed over the Bloch sphere. We guess at random that the state is $|\phi\rangle$. On the average, what is the *fidelity* F of our guess, defined by

$$F \equiv |\langle\phi|\psi\rangle|^2 \ . \tag{2.122}$$

2.2 Fidelity after measurement

After randomly selecting a one-qubit pure state as in the previous problem, we perform a measurement of the spin along the $\hat{z}$-axis. This measurement prepares a state described by the density matrix

$$\boldsymbol{\rho} = \mathbf{P}_\uparrow\langle\psi|\mathbf{P}_\uparrow|\psi\rangle + \mathbf{P}_\downarrow\langle\psi|\mathbf{P}_\downarrow|\psi\rangle \tag{2.123}$$

(where $\mathbf{P}_{\uparrow,\downarrow}$ denote the projections onto the spin-up and spin-down states along the $\hat{z}$-axis). On the average, with what fidelity

$$F \equiv \langle\psi|\boldsymbol{\rho}|\psi\rangle \tag{2.124}$$

does this density matrix represent the initial state $|\psi\rangle$? (The improvement in F compared to the answer to the previous problem is a crude measure of how much we learned by making the measurement.)

2.3 Schmidt decomposition

For the two-qubit state

$$\Phi = \frac{1}{\sqrt{2}}|\uparrow\rangle_A\left(\frac{1}{2}|\uparrow\rangle_B + \frac{\sqrt{3}}{2}|\downarrow\rangle_B\right) + \frac{1}{\sqrt{2}}|\downarrow\rangle_A\left(\frac{\sqrt{3}}{2}|\uparrow\rangle_B + \frac{1}{2}|\downarrow\rangle_B\right), \tag{2.125}$$

a. Compute $\boldsymbol{\rho}_A = \mathrm{tr}_B\left(|\Phi\rangle\langle\Phi|\right)$ and $\boldsymbol{\rho}_B = \mathrm{tr}_A\left(|\Phi\rangle\langle\Phi|\right)$.

b. Find the Schmidt decomposition of $|\Phi\rangle$.

2.4 Tripartite pure state

Is there a Schmidt decomposition for an arbitrary *tripartite* pure state? That is if $|\psi\rangle_{ABC}$ is an arbitrary vector in $\mathcal{H}_A\otimes\mathcal{H}_B\otimes\mathcal{H}_C$, can we find orthonormal bases $\{|i\rangle_A\}$, $\{|i\rangle_B\}$, $\{|i\rangle_C\}$ such that

$$|\psi\rangle_{ABC} = \sum_i \sqrt{p_i}|i\rangle_A\otimes|i\rangle_B\otimes|i\rangle_C \ ? \tag{2.126}$$

Explain your answer.

2.5 Quantum correlations in a mixed state

Consider a density matrix for two qubits

$$\boldsymbol{\rho} = \frac{1}{8}\mathbf{1} + \frac{1}{2}|\psi^-\rangle\langle\psi^-| \ , \tag{2.127}$$

where $\mathbf{1}$ denotes the 4×4 unit matrix, and

$$|\psi^-\rangle = \frac{1}{\sqrt{2}}\left(|\uparrow\rangle|\downarrow\rangle - |\downarrow\rangle|\uparrow\rangle\right) \ . \tag{2.128}$$

Suppose we measure the first spin along the $\hat{n}$ axis and the second spin along the $\hat{m}$ axis, where $\hat{n}\cdot\hat{m} = \cos\theta$. What is the probability that both spins are "spin-up" along their respective axes?