Chapter 3: Dynamics, Pulses, and Steering Quantum Systems

Quantum control begins when we stop asking only, “How does this quantum system naturally evolve?” and begin asking, “How can we change that evolution on purpose?”

In Chapter 2, we learned that a closed quantum system evolves according to the Schrödinger equation. The central object controlling that evolution is the Hamiltonian, the operator associated with the system’s energy. In this chapter, we make the next step: we allow the Hamiltonian to depend on time because an experimentalist is applying a laser, microwave field, radio-frequency field, voltage pulse, magnetic-field pulse, or another external influence.

That is the basic idea of steering:

$$\text{natural dynamics} + \text{controlled time-dependent influence} = \text{designed quantum mot}$$

This chapter focuses on the simplest and most important model in quantum control: a two-level system. A two-level system has two relevant quantum states, often written as $|0\rangle$ and $|1\rangle$, or as a ground state $|g\rangle$ and excited state $|e\rangle$. This model describes many real systems approximately: atomic transitions, electron spins, nuclear spins, superconducting qubits, trapped-ion qubits, and many optical transitions. The two-level model is not the whole of quantum control, but it is the best place to learn the language of pulses, resonance, phase, and rotations (Allen and Eberly, 1987; Levitt, 2008).

By the end of this chapter, you should understand how a simple pulse can turn $|0\rangle$ into $|1\rangle$, how a shorter pulse can create a superposition, why resonance matters, and why the phase of a pulse is not a small detail but a genuine control knob.

3.1 The Hamiltonian as the engine of motion

The Schrödinger equation for a pure state is

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

Here:

- |ψ(t)⟩ is the state at time t,
- H(t) is the Hamiltonian, possibly changing in time,
- hbar is the reduced Planck constant.

If the Hamiltonian is time independent, the evolution is generated by

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

so that

$$|\psi(t)\rangle = U(t)|\psi(0)\rangle.$$

The operator U(t) is called a unitary time-evolution operator. “Unitary” means that it preserves the length of the quantum state vector, which is necessary because total probability must remain equal to 1 in a closed system. This is standard quantum mechanics: for a closed system, the Schrödinger equation gives deterministic, norm-preserving evolution until measurement is performed (Sakurai and Napolitano, 2020).

In control, the Hamiltonian is usually split into two parts:

$$H(t) = H_0 + H_c(t).$$

The first part, H_0, is the drift Hamiltonian. It describes what the system does naturally when we do not actively control it. The second part, $H_{(c)}(t)$, is the control Hamiltonian. It describes the effect of external fields that we choose.

For example:

- For an atom, H_0 describes the atom’s internal energy levels, while $H_{(c)}(t)$ may describe interaction with a laser field.
- For a spin, H_0 may describe precession in a static magnetic field, while $H_{(c)}(t)$ may describe a radio-frequency or microwave pulse.
- For a superconducting qubit, H_0 describes the qubit’s natural transition frequency, while $H_{(c)}(t)$ may describe a microwave voltage or current pulse.

This decomposition is simple, but it is one of the most useful ideas in quantum control:

The system drifts on its own, and we steer it by adding controlled terms.

3.2 A two-level system: the basic control target

Let us begin with two states, $|0\rangle$ and $|1\rangle$. A general pure state is

$$|\psi\rangle = \alpha|0\rangle + \beta|1\rangle,$$

where α and β are complex amplitudes satisfying

$$|\alpha|^2 + |\beta|^2 = 1.$$

The probability of measuring $|0\rangle$ is $|\alpha|^2$, and the probability of measuring $|1\rangle$ is $|\beta|^2$.

Suppose the two states have an energy difference

$$E_1 - E_0 = \hbar\omega_0.$$

The angular frequency ω_0 is the system's transition frequency. It tells us how quickly the relative phase between $|0\rangle$ and $|1\rangle$ changes under free evolution.

A useful model for the uncontrolled Hamiltonian is

$$H_0 = \frac{\hbar\omega_0}{2}\sigma_z,$$

up to a choice of energy zero and sign convention. Here $\sigma_{(z)}$ is one of the Pauli operators. The Pauli operators σ_x, $\sigma_{(y)}$, and $\sigma_{(z)}$ are basic 2×2 matrices that generate rotations of a two-level quantum state on the Bloch sphere. They are central in both spin physics and quantum information (Nielsen and Chuang, 2010; Levitt, 2008).

The exact sign convention is not important for the main physical idea. What matters is this:

> The two-level system has a natural frequency ω_0, and an external field can drive transitions most effectively when it oscillates near that frequency.

That idea is called resonance.

3.3 External fields: how the laboratory enters the Hamiltonian

A quantum system is usually controlled through its interaction with an external field.

For an atom or molecule, an electric field E(t) can interact with an electric dipole moment. In simplified form, the interaction energy is often written as

$$H_c(t) = -\mathbf{d} \cdot \mathbf{E}(t),$$

where d is the electric dipole operator and E(t) is the applied electric field. This is a standard starting point for light-matter interaction in atomic and molecular quantum mechanics (Cohen-Tannoudji, Diu, and Laloë, 1977).

For a spin in a magnetic field, the interaction has the form

$$H_c(t) = -\boldsymbol{\mu} \cdot \mathbf{B}(t),$$

where boldsymbolμ is the magnetic moment and B(t) is the magnetic field. This is the basic interaction used in nuclear magnetic resonance, electron spin resonance, and many spin-qubit control methods (Levitt, 2008).

In both cases, the control field is a function of time. A typical driving field may look like

$$E(t) = A(t)\cos(\omega_d t + \phi).$$

This expression contains three important control knobs:

1. Amplitude, A(t): how strong the field is.
2. Drive frequency, $\omega_{(d)}$: how fast the field oscillates.
3. Phase, φ: where the oscillation starts in its cycle.

The function A(t) is often called the envelope of the pulse. The cosine oscillation is the fast carrier wave, and the envelope tells us how the pulse turns on, changes strength, and turns off.

For example, a microwave pulse applied to a qubit might have a carrier frequency near the qubit transition frequency, while its envelope might be a square, Gaussian, or smoothly shaped function. The carrier frequency decides which transition is addressed; the envelope and duration decide how much the state is rotated.

3.4 Time-dependent Hamiltonians and why order matters

When H(t) changes in time, the evolution is more subtle than simply writing e^(-iHt/hbar), because there is no single fixed Hamiltonian.

If the Hamiltonian at different times commutes with itself,

$$[H(t_1), H(t_2)] = 0$$

for all times t_1 and t_2, then the evolution can be written as

$$U(t) = \exp\left[-\frac{i}{\hbar}\int_0^t H(t')\,dt'\right].$$

But in quantum control, Hamiltonians at different times often do not commute. That means the order of operations matters.

This is not only a mathematical issue. It has a physical meaning. A rotation about the x-axis followed by a rotation about the y-axis is generally not the same as doing the y-rotation first and then the x-rotation. Quantum control inherits this noncommutativity from the operator structure of quantum mechanics.

For a general time-dependent Hamiltonian, the evolution is written using a time-ordered exponential:

$$U(t) = \mathcal{T}\exp\left[-\frac{i}{\hbar}\int_0^t H(t')\,dt'\right].$$

The symbol $\mathcal{T}$ means “time ordering.” It tells us that earlier and later Hamiltonians must be applied in the correct chronological order. This formal solution is standard in time-dependent quantum mechanics (Sakurai and Napolitano, 2020).

At first, this may look abstract. But the lesson is practical:

> A pulse sequence is not just a list of pulse areas. The order, timing, phase, and frequency of the pulses all matter.

3.5 Resonance: pushing at the right frequency

Resonance is one of the most important ideas in all of control physics. A child on a swing moves higher when you push at the right rhythm. A radio receives one station when its circuit is tuned to the correct frequency. A two-level quantum system responds strongly when the applied field oscillates near its transition frequency.

Suppose a two-level system has transition frequency ω_0. We apply a drive at frequency $\omega_{(d)}$. The difference

$$\Delta = \omega_0 - \omega_d$$

is called the detuning.

- If $\Delta=0$, the drive is on resonance.
- If $|\Delta|$ is small, the drive is near resonance.
- If $|\Delta|$ is large, the drive is far off resonance.

On resonance, the drive efficiently transfers population between the two states. Off resonance, the transfer is less efficient, and the state may mostly acquire phase rather than fully switch states.

This is a central reason why quantum control can be selective. If a system has several transitions at different frequencies, a drive can be tuned near one transition while mostly avoiding others. This frequency selectivity is widely used in atomic physics, magnetic resonance, and qubit control (Allen and Eberly, 1987; Levitt, 2008).

3.6 Rabi oscillations: the simplest controlled quantum motion

Now we reach one of the fundamental phenomena of driven two-level systems: Rabi oscillation.

A Rabi oscillation is the periodic transfer of population between two quantum states caused by a near-resonant external drive. The effect is named after Isidor I. Rabi, whose work on magnetic resonance helped establish how oscillating fields can drive transitions between quantum states. In modern quantum control, Rabi oscillations are a basic calibration tool: they tell us how strongly a pulse couples to a qubit, atom, ion, molecule, or spin.

In a convenient rotating-frame description, which we will explain shortly, a resonantly driven two-level system can be represented by the Hamiltonian

$$H_{\text{eff}} = \frac{\hbar\Omega}{2}\sigma_x.$$

Here Ω is the Rabi frequency. It measures how fast the drive rotates the quantum state. If the system starts in |0⟩, then under this Hamiltonian the probability of finding it in |1⟩ is

$$P_1(t) = \sin^2\left(\frac{\Omega t}{2}\right).$$

This equation has a simple interpretation:

- At $t=0$, $P_1(0)=0$. The system is still in |0⟩.
- At $t=\pi/\Omega$, $P_1=1$. The system has been transferred to |1⟩.
- At $t=2\pi/\Omega$, $P_1=0$. The system has returned to |0⟩.

So the drive causes the population to slosh back and forth between the two states.

A pulse of duration

$$t_\pi = \frac{\pi}{\Omega}$$

is called a π-pulse. It rotates the state by angle π on the Bloch sphere and transfers |0⟩ to |1⟩, up to an overall global phase. A pulse of duration

$$t_{\pi/2} = \frac{\pi}{2\Omega}$$

is called a π/2-pulse. It creates an equal superposition of |0⟩ and |1⟩, with a relative phase determined by the pulse phase.

Rabi oscillations and pulse rotations are standard tools in optical resonance, magnetic resonance, and qubit control (Allen and Eberly, 1987; Levitt, 2008).

3.7 What detuning does to Rabi oscillations

If the drive is not exactly resonant, the effective Hamiltonian becomes

$$H_{\text{eff}} = \frac{\hbar}{2}\left(\Delta\sigma_z + \Omega\sigma_x\right),$$

where $\Delta=\omega_0-\omega_{(d)}$ is the detuning and Ω is the resonant Rabi frequency.

The system now rotates about a tilted axis on the Bloch sphere. The effective rotation frequency is

$$\Omega_{\text{eff}} = \sqrt{\Omega^2 + \Delta^2}.$$

If the system starts in $|0\rangle$, the transition probability becomes

$$P_1(t) = \frac{\Omega^2}{\Omega^2 + \Delta^2}\sin^2\left(\frac{\sqrt{\Omega^2 + \Delta^2}\,t}{2}\right).$$

This formula shows two important effects.

First, detuning changes the oscillation frequency from Ω to $\Omega_{(\text{eff})}$. Second, detuning reduces the maximum possible population transfer. The prefactor

$$\frac{\Omega^2}{\Omega^2 + \Delta^2}$$

is less than 1 whenever $\Delta \neq 0$. Therefore, a strongly detuned drive cannot fully invert the system, no matter how long we wait, in this simple constant-drive model.

This is why calibration matters. If your pulse frequency is slightly wrong, a pulse designed to be a perfect π-pulse may fail. In a quantum computer, that means a gate error. In a quantum sensor, that means reduced contrast. In spectroscopy, that means a distorted measurement.

3.8 The rotating frame: making fast motion understandable

The laboratory frame description of a driven two-level system can be visually confusing. The applied field may oscillate billions or trillions of times per second, while the useful controlled motion of the quantum state may be much slower. To see the control clearly, physicists often use a rotating frame.

A rotating frame is a coordinate system that rotates at or near the drive frequency. Instead of watching the state from a fixed laboratory viewpoint, we watch it from a viewpoint that spins along with the applied field.

This is similar to watching a runner from a moving vehicle. If you stand on the sidewalk, the runner moves quickly past you. If you ride beside the runner at nearly the same speed, the runner's motion relative to you looks slow. The rotating frame removes much of the rapid oscillation and reveals the slower dynamics that the pulse actually controls.

Mathematically, changing to a rotating frame means applying a time-dependent unitary transformation to the state and Hamiltonian. If the transformation is R(t), then the state in the rotating frame is related to the laboratory-frame state by a rule of the form

$$|\psi_{\mathrm{rot}}(t)\rangle = R^{\dagger}(t)|\psi_{\mathrm{lab}}(t)\rangle.$$

The Hamiltonian also changes. The rotating-frame Hamiltonian contains both the transformed original Hamiltonian and an extra term due to the moving frame itself. This is a standard feature of time-dependent changes of picture in quantum mechanics (Sakurai and Napolitano, 2020).

For a near-resonant driven two-level system, after moving to the rotating frame and making a common approximation called the rotating-wave approximation, the Hamiltonian often takes the simple form

$$H_{\mathrm{rot}} = \frac{\hbar}{2}\left[\Delta\sigma_z + \Omega(t)\left(\cos\phi\,\sigma_x + \sin\phi\,\sigma_y\right)\right].$$

This compact equation is one of the most important equations in elementary quantum control.

Let us unpack it carefully.

- Δ is the detuning. It produces rotation about the z-axis.

- Ω(t) is the time-dependent Rabi frequency. It sets the strength of the control.

- φ is the phase of the drive.

- cosφ σ_x+sinφ $\sigma_{(y)}$ means that the pulse rotates the state about an axis in the xy-plane of the Bloch sphere.

The rotating-wave approximation removes rapidly oscillating terms that average out when the drive is near resonance and not too strong compared with the transition frequency. This approximation is widely used in magnetic resonance, atomic physics, and quantum optics, but it has limits: if the drive is extremely strong, very short, or far from resonance, the neglected fast terms may become important (Allen and Eberly, 1987; Levitt, 2008).

3.9 Pulse area: how much rotation a pulse produces

On resonance, where Δ=0, and for a fixed phase φ, the rotating-frame Hamiltonian is

$$H_{\mathrm{rot}} = \frac{\hbar\Omega(t)}{2}\left(\cos\phi\,\sigma_x + \sin\phi\,\sigma_y\right).$$

If the phase stays fixed, the rotation axis stays fixed. The amount of rotation is determined by the pulse area

$$\theta = \int_{t_i}^{t_f} \Omega(t)\,dt.$$

The angle θ is measured in radians. It tells us how far the state rotates on the Bloch sphere.

For a square pulse with constant Ω and duration T,

$$\theta = \Omega T.$$

So:

- A π-pulse satisfies Ω T=π.

- A π/2-pulse satisfies Ω T=π/2.

- A 2π-pulse satisfies $\Omega T = 2\pi$.

For a shaped pulse, $\Omega(t)$ changes over time, but the pulse area still gives the rotation angle as long as the pulse remains resonant, the phase is fixed, and the effective Hamiltonian points along the same Bloch-sphere axis throughout the pulse.

This is a powerful simplification. It means that a short strong pulse and a longer weak pulse can produce the same ideal rotation if their pulse areas are equal. For example, if one pulse has twice the Rabi frequency but half the duration, the product ΩT is the same, so the ideal rotation angle is the same.

However, in real systems, “same area” does not always mean “same performance.” A very short pulse may have broad frequency content and accidentally affect nearby transitions. A very long pulse may allow decoherence to damage the state. This is one reason practical pulse design is a balance between speed, selectivity, and robustness.

3.10 Phase control: choosing the rotation axis

The pulse area determines how far the state rotates. The pulse phase determines the direction of the rotation axis in the equatorial plane of the Bloch sphere.

Recall the resonant rotating-frame Hamiltonian:

$$H_{\text{rot}} = \frac{\hbar\Omega(t)}{2}\left(\cos\phi\,\sigma_x + \sin\phi\,\sigma_y\right).$$

If $\varphi=0$, then

$$H_{\text{rot}} = \frac{\hbar\Omega(t)}{2}\sigma_x.$$

The pulse produces an x-axis rotation.

If $\varphi=(\pi)/(2)$, then

$$H_{\text{rot}} = \frac{\hbar\Omega(t)}{2}\sigma_y.$$

The pulse produces a y-axis rotation.

Intermediate phases produce rotations about axes between x and y.

This matters because quantum amplitudes are complex. Creating a superposition is not only about the probabilities $|\alpha|^2$ and $|\beta|^2$. It is also about the relative phase between α and β. Two states can have the same measurement probabilities in the $|0\rangle,|1\rangle$ basis but behave differently in later interference experiments because their relative phases differ.

For example, starting from $|0\rangle$, an ideal resonant x-axis π/2-pulse produces a state of the form

$$\frac{1}{\sqrt{2}}\left(|0\rangle - i|1\rangle\right),$$

up to convention-dependent signs. A y-axis π/2-pulse can produce

$$\frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right).$$

Both are equal superpositions in the sense that measuring in the $|0\rangle,|1\rangle$ basis gives probability 1/2 for each outcome. But they are not the same quantum state, because the relative phase between $|0\rangle$ and $|1\rangle$ is different.

This is why phase control is essential in quantum computing and spectroscopy. Quantum gates are not only population transfers; they are controlled transformations of amplitudes and phases (Nielsen and Chuang, 2010; Levitt, 2008).

3.11 Simple pulses as quantum gates

In quantum computing, a one-qubit gate is a unitary operation on a two-level system. The same mathematics appears in many other quantum-control settings, even when we do not call the operation a “gate.”

A resonant pulse with area θ and phase φ produces the unitary operation

$$U(\theta,\phi) = \exp\left[-\frac{i\theta}{2}\left(\cos\phi\,\sigma_x + \sin\phi\,\sigma_y\right)\right].$$

This is a rotation by angle θ about an axis in the xy-plane.

Some important examples are:

$$X_\pi = e^{-i\pi\sigma_x/2},$$

which flips |0⟩ and |1⟩ up to phases. This is often called an X gate or bit-flip operation.

A π/2-rotation about x is

$$X_{\pi/2} = e^{-i(\pi/2)\sigma_x/2}.$$

It creates a superposition from a basis state.

A π/2-rotation about y is

$$Y_{\pi/2} = e^{-i(\pi/2)\sigma_y/2}.$$

It creates a superposition with a different relative phase.

These rotations are the building blocks of more complicated control sequences. In a quantum computer, they are single-qubit gates. In magnetic resonance, they are pulse rotations. In atomic physics, they are coherent state manipulations. The physical platforms differ, but the two-level control mathematics is shared.

This shared mathematics is one reason quantum control is such a unifying field. The same Bloch-sphere picture can describe a nuclear spin in a magnetic field, an electron spin in a defect center, an optical transition in an atom, or a superconducting qubit driven by microwaves.

3.12 A worked example: making a superposition

Let us design a simple control task.

Goal: Start in |0⟩ and create

$$|\psi_{\text{target}}\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right).$$

This is an equal superposition with zero relative phase.

Assume we can apply a resonant pulse with phase $\varphi=\pi/2$. That means the Hamiltonian is proportional to $\sigma_{(y)}$:

$$H_{\text{rot}} = \frac{\hbar\Omega}{2}\sigma_y.$$

The corresponding unitary is

$$U_y(\theta) = e^{-i\theta\sigma_y/2}.$$

If we choose $\theta=\pi/2$, then

$$U_y(\pi/2)|0\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right),$$

up to the basis and sign convention used for $\sigma_{(y)}$.

The required pulse duration is

$$T = \frac{\theta}{\Omega} = \frac{\pi}{2\Omega}.$$

So the procedure is:

1. Tune the drive to resonance: $\omega_{(d)}=\omega_0$.
2. Choose the pulse phase to select the y-axis.
3. Choose the pulse duration so that Ω $T=\pi/2$.
4. Turn off the pulse.

In an ideal closed system, this produces the target superposition. In a real laboratory, we would then measure the result many times to estimate whether the desired probabilities and phases were actually achieved.

3.13 Another worked example: population inversion

Now suppose the goal is simpler.

Goal: Start in $|0\rangle$ and transfer all population to $|1\rangle$.

This is called population inversion. In a two-level system, it means changing the state from one basis state to the other.

Use a resonant pulse with $\varphi=0$, so the Hamiltonian is

$$H_{\mathrm{rot}} = \frac{\hbar\Omega}{2}\sigma_x.$$

The transition probability is

$$P_1(t) = \sin^2\left(\frac{\Omega t}{2}\right).$$

To make $P_1=1$, we need

$$\frac{\Omega T}{2} = \frac{\pi}{2},$$

so

$$T = \frac{\pi}{\Omega}.$$

This is a π-pulse.

In ideal theory,

$$|0\rangle \longrightarrow |1\rangle$$

up to an overall phase. Overall global phases do not affect measurement probabilities, but relative phases between components of a superposition do matter.

This simple operation appears everywhere:

- flipping a spin,
- exciting an atom,
- applying a single-qubit X gate,
- preparing an excited state for spectroscopy,
- transferring population between two selected molecular states.

The same mathematics is used across many physical systems, even though the hardware may be very different.

3.14 Pulse sequences: steering by composing rotations

A single pulse can do a lot, but many control tasks require a sequence of pulses. A pulse sequence is an ordered set of pulses with specified phases, durations, amplitudes, and delays.

For example, consider the sequence

$$X_{\pi/2} \quad \text{then} \quad Y_{\pi/2}.$$

This means we first rotate the state by π/2 about the x-axis, then rotate it by π/2 about the y-axis.

Because rotations about different axes do not generally commute,

$$X_{\pi/2}Y_{\pi/2} \neq Y_{\pi/2}X_{\pi/2}.$$

The order matters.

This noncommutativity is not a nuisance. It is a resource. By composing rotations about different axes, we can build many different final operations. In fact, for a single qubit, rotations about two nonparallel axes are enough to construct arbitrary single-qubit rotations, up to practical limitations such as calibration errors and decoherence. This fact is central in quantum information theory and in experimental pulse design (Nielsen and Chuang, 2010; Levitt, 2008).

A useful way to think about pulse sequences is this:

> Each pulse is a controlled rotation. A pulse sequence is a route through state space.

The route matters. If the system is noisy, some routes may be faster or more robust than others. If the hardware has limited amplitude or bandwidth, some routes may be possible while others are not. Later, in optimal control, we will learn how computers can search for high-performing routes automatically.

3.15 Free evolution as a kind of control

So far, we have emphasized applied pulses. But sometimes doing nothing for a chosen time is also part of control.

During a delay, the system evolves under its drift Hamiltonian H_0. For a two-level system, this often corresponds to a rotation about the z-axis of the Bloch sphere. If the state is a superposition,

$$|\psi\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + |1\rangle\right),$$

then free evolution changes the relative phase between $|0\rangle$ and $|1\rangle$. After time t, the state may become, up to an overall phase,

$$|\psi(t)\rangle = \frac{1}{\sqrt{2}}\left(|0\rangle + e^{-i\omega_0 t}|1\rangle\right),$$

depending on the sign convention used for the Hamiltonian.

The measurement probabilities in the

Document information

Chapter 3: Dynamics, Pulses, and Steering Quantum Systems

Project	Quantum Control in Action
Document	Document 1.7
Author	mujirin
Verifier	Not verified
Downloaded	July 04, 2026 22:23 KST
Status	Working
Document link	https://www.theorytrace.com/projects/quantum-control-in-action/documents/chapter-3--dynamics-pulses-and-steering-quantum-systems/