Elementary Particle Physics

Notes

Abhishek Bohare

a_bohare@icloud.com

typeset with LaTeX

Contents

1 The Four Fundamental Forces **1**

1.1 Introduction: four forces, one classification 1

1.1.1 What “force” means here 1

1.1.2 Three questions to ask of any force 1

1.2 The gravitational force 3

1.2.1 Newton’s law, and why it explains so much 4

1.2.2 A worked example: how weak is gravity, really? 4

1.2.3 Where Newton’s law breaks down, and Einstein’s alternative 5

1.3 The electromagnetic force 7

1.3.1 Classical description: Coulomb and Maxwell 7

1.3.2 The quantum picture: QED 7

1.3.3 Reading a Feynman diagram: three simple rules 8

1.4 The strong nuclear force 9

1.4.1 The puzzle the strong force has to solve 9

1.4.2 Quarks are not optional: what is inside a proton 10

1.4.3 Why colour must be “colourless” overall 10

1.4.4 Confinement: why free quarks are never seen 10

1.4.5 The residual strong force, and why the nucleus has a finite size 11

1.5 The weak nuclear force 12

1.5.1 Why “weak” is a relative, not an absolute, term 12

1.5.2 Why a massive carrier makes a force look weak 12

1.5.3 What makes the weak force unique: changing flavour 13

1.5.4 Worked example: beta decay, at three levels of detail 13

1.5.5 Electroweak unification 14

1.5.6 Why stars need the weak force 15

1.5.7 Parity violation: a preview 15

1.6 Summary and comparison 15

1.6.1 The comparison table 15

1.6.2 A diagnostic habit worth building 16

2 Antiparticles and Dirac Hole Theory **18**

2.1 Antiparticles: the basic idea 18

2.1.1 Annihilation and pair production 19

2.2 Dirac hole theory 20

2.2.1 The problem: negative energy 20

2.2.2 The Dirac sea 20

2.2.3 Why hole theory was eventually superseded 22

2.3 Crossing symmetry 23

2.3.1 The rule, stated precisely 23

2.3.2 Worked example 23

3 Neutrinos **25**

3.1 The neutrino hypothesis 25

3.1.1 A crisis in beta decay 25

3.1.2 Pauli's "desperate remedy" 26

3.1.3 Fermi's theory, and getting the name right 26

3.2 Properties of the neutrino 26

3.2.1 No electric charge 27

3.2.2 A mass far smaller than any charged lepton's, but (as it turns out) not quite zero 27

3.2.3 Spin $\frac{1}{2}$, and three flavours 27

3.2.4 A particle defined almost entirely by what it does *not* do . 27

3.3 Handedness and parity 28

3.3.1 Helicity: spin relative to motion 28

3.3.2 An empirical rule of the weak interaction 29

3.3.3 The link to parity 29

4 Parity Violation in Beta Decay: the Wu Experiment **31**

4.1 Parity as a symmetry 31

4.1.1 What a conserved symmetry actually buys you 31

4.1.2 Defining parity concretely 31

4.1.3 Why this went untested for so long 31

4.2 The Wu experiment . . . 32
4.2.1 Setting up the test . . . 32
4.2.2 The logic of the test, spelled out in full . . . 33
4.2.3 The result . . . 33
4.2.4 Connecting the asymmetry to Chapter 3's handedness rule 34
4.2.5 The aftermath . . . 34

5 Neutrino Oscillations and the Solar Neutrino Problem 35
5.1 The solar neutrino problem . . . 35
5.1.1 Why the Sun should produce a precisely predictable flux of neutrinos . . . 35
5.1.2 The experiments, and a persistent shortfall . . . 35
5.2 Neutrino oscillations . . . 36
5.2.1 Flavour states versus mass states . . . 36
5.2.2 Why propagation makes the mixture change . . . 37
5.2.3 How this resolves the solar neutrino problem . . . 37

6 Leptons 39
6.1 Classifying the particle zoo . . . 39
6.1.1 A history of classification, not just a list . . . 39
6.2 The lepton family . . . 39
6.2.1 Stability . . . 41
6.3 Why leptons are leptons . . . 41
6.3.1 The real reason, stated precisely . . . 41
6.3.2 A worked check: can a lepton decay into a hadron? . . . 42
6.3.3 Lepton number, previewed . . . 42

7 Hadrons 43
7.1 Hadrons: particles built from quarks . . . 43
7.1.1 The defining criterion . . . 43
7.1.2 Two families, from one simple rule . . . 43
7.2 Pions: the lightest mesons . . . 44
7.3 Some familiar baryons . . . 45
7.3.1 Why the neutron is stable inside a nucleus but not outside it 45

7.3.2 From individual particles to a whole spectrum 45

8 Lepton Number, Baryon Number and Strangeness 47

8.1 Lepton number and baryon number 47

8.1.1 Why bookkeeping rules like this are needed at all 47

8.1.2 Lepton number 47

8.1.3 Baryon number 48

8.2 Strangeness 48

8.2.1 The puzzle that gave the quantum number its name 48

8.2.2 Gell-Mann and Nishijima's resolution 49

8.2.3 A general pattern, not a one-off curiosity 50

9 The Eightfold Way 51

9.1 Organising the hadrons 51

9.1.1 Following the pattern that worked before 51

9.1.2 The baryon octet, built up piece by piece 51

9.1.3 What the pattern was telling everyone 52

9.2 Prediction and discovery of the Ω^- 52

9.2.1 A second pattern: the decuplet 52

9.2.2 The prediction, and why it was so bold 53

9.2.3 The discovery, and what it proved 53

10 Quarks: Colour and Flavour 54

10.1 The quark model 54

10.1.1 From pattern to substance 54

10.1.2 Six flavours, discovered one at a time 54

10.1.3 Why only u and d matter for ordinary matter 55

10.1.4 A genuinely predictive accounting scheme 55

10.2 Why colour was needed 56

10.2.1 A problem the quark model had from day one 56

10.2.2 The resolution: a new quantum number 56

10.2.3 Confinement, revisited with colour explicit 57

11 Tetraquarks and Pentaquarks 58

11.1 Exotic hadrons 58
11.1.1 Colour neutrality is more permissive than it first looks . . . 58
11.1.2 Why “theoretically allowed” took fifty years to become “experimentally confirmed” 58
11.2 Confirmed discoveries at LHCb 59
11.2.1 The experiment that finally settled the question 59
11.2.2 Two competing pictures 59
11.2.3 Why this chapter closes the course where it does 60

Todo list

- The lecture's closing chapter ("Summary and comparison", from 40:02) recaps points already covered in Sections 1.2–1.5 above rather than introducing new material; the comparison table, bar chart and the diagnostic checklist above are my own synthesis and extension of the figures and properties given earlier in the lecture, cross-checked against the standard values quoted for each force. 17

1 The Four Fundamental Forces

1.1 Introduction: four forces, one classification

Why should there be exactly four forces, rather than one, or fifty? The honest answer is that nobody derived the number four from first principles: it is an empirical fact about this universe, established by more than a century of experiments that kept turning up the same four, qualitatively distinct, patterns of interaction, and no fifth. The task of this chapter is to make that classification concrete enough to use, following the first lecture of the *Elementary Particle Physics* series (For the Love of Physics, YouTube), which surveys all four before the series moves on to the particles themselves.

1.1.1 What "force" means here

Before listing the four, let's be precise about what counts as one, because everyday language is not much help. In everyday speech "force" can mean almost any push or pull: friction, the normal force from a table, the tension in a rope, air resistance. None of these is fundamental. Friction, for instance, is nothing but the electromagnetic force acting between the atoms of two surfaces in contact, mediated by the same electron clouds that hold each surface's atoms together; there is no separate "friction force" at the level of fundamental physics, only electromagnetism expressing itself in a complicated, many-body way. A **fundamental** force, by contrast, is one that cannot be explained as some other force operating in disguise. By that standard, every interaction we have ever observed, from the tension in a rope to the binding of a nucleus, reduces to exactly one of four: **gravity**, the **electromagnetic force**, the **strong interaction** and the **weak interaction**.

1.1.2 Three questions to ask of any force

A fast way to tell the four forces apart, and the organising structure for the rest of this chapter, is to ask the same three questions of each one.

1. **What does it couple to?** Every force acts only on particles carrying some specific property: gravity couples to mass-energy, electromagnetism to electric charge, the strong force to colour charge, the weak force to flavour (and, like the others, to the weak charge each particle carries). A particle with zero of the relevant charge is simply invisible to that force. A neutrino, carrying no electric charge and no colour, is blind to electromagnetism and the strong force; it only feels gravity and the weak force.
2. **How far does it reach?** Some forces have infinite range (gravity, electromagnetism); others are confined to subatomic distances (the strong and weak forces). This is directly tied to whether the force's carrier particle is massless

(giving infinite range, as for the photon) or massive (giving a short range, as for the W and Z bosons of the weak force, by an argument given in Section 1.5).

3. **How strong is it, at a given distance, compared with the others?** This only makes sense once a common distance scale is fixed, since the forces have such different ranges; the conventional choice is to compare them all at the scale of a nucleus, roughly a femtometre, where all four can act simultaneously and be measured against one another. Section 1.6 collects these comparisons.

Keep these three questions in your back pocket; they are the fastest route through the rest of the chapter, and the table in Section 1.6 is exactly this information collected in one place.

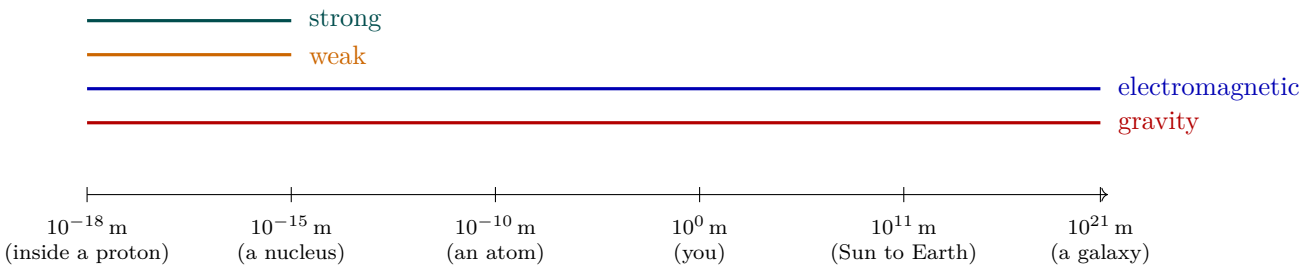


Figure 1.1: How far each force actually reaches, drawn on a logarithmic scale running from inside a proton to the size of a galaxy. The strong and weak forces (top two bars) simply stop, abruptly, at nuclear distances; gravity and electromagnetism (bottom two bars) keep going, in principle, forever. Come back to this picture after each of Sections 1.2–1.5: every force in this chapter belongs on one of these four lines.

A preview of the pattern

Every one of the four forces, in the modern (standard-model) account, arises from the exchange of a specific **carrier particle** between the particles that are interacting: the graviton for gravity, the photon for electromagnetism, gluons for the strong force, and the $W^{\pm}/Z^{0}$ bosons for the weak force. This exchange-particle picture, explained properly in Section 1.3, is the single idea that unifies the treatment of all four forces. Hold onto it as a thread through the whole chapter.
Terms used here, defined for later reference:

- **Carrier particle** (or **mediator**, or **gauge boson**): the particle exchanged between two interacting particles that *is* the force, in the sense that the attraction or repulsion is modelled as a continuous back-and-forth exchange of these particles, each exchange transferring a little momentum. All carrier particles are **bosons** (integer spin, specifically spin 1 for every carrier particle listed here), as opposed to the **fermions** (half-integer spin, $\frac{1}{2}, \frac{3}{2}, \ldots$) that make up matter itself (electrons, quarks, neutrinos).
- **Graviton**: the hypothesised carrier of gravity, predicted to be massless and to have spin 2; unlike the other three carriers, it has never been experimentally detected, and no full, mathematically consistent quantum theory of it yet exists, which is part of why gravity is normally treated separately from the other three forces (together called the **electroweak and strong**, or loosely the "particle-physics", forces) in the standard model.
- **Photon** (γ): the carrier of the electromagnetic force; massless, spin 1, electrically neutral. Covered in depth in Section 1.3.
- **Gluon** (g): the carrier of the strong force; massless, spin 1, but unlike the photon it carries colour charge itself (Section 1.4), which lets gluons interact with *each other*, not only with quarks, a key difference from electromagnetism explored in Section 1.4.
- $W^{\pm}$ **and** Z^{0} **bosons**: the two carriers of the weak force; the $W^{\pm}$ bosons are electrically charged ($+e$ or $-e$) and the Z^{0} is neutral. All three are massive (around 80–91 GeV/c^{2}, Section 1.5), which is the single fact responsible for the weak force's short range.
- **Standard model**: the overall theoretical framework, built up through the twentieth century and essentially complete by the 1970s, describing the electromagnetic, weak and strong forces and every known matter particle (quarks and leptons) within one consistent quantum field theory. It deliberately excludes gravity.

1.2 The gravitational force

Gravity is the force everyone meets first, and also the force that is easiest to get wrong about at the subatomic scale. It governs essentially all macroscopic and astronomical structure, why we stay on the Earth's surface, why the planets orbit

the Sun, why galaxies hold together, why a collapsing cloud of gas can become hot enough at its core to ignite fusion and become a star, and yet, as will become clear by the end of this section, it is utterly negligible between two individual particles such as two protons. Both of these statements are true simultaneously, and understanding why is the first real lesson of this chapter.

1.2.1 Newton's law, and why it explains so much

The first quantitative description of gravity is Newton's law of universal gravitation: two point masses m_1 and m_2, separated by a distance r, attract one another with a force

$$F = \frac{G m_1 m_2}{r^2}, \tag{1.2.1}$$

where $G = 6.674 \times 10^{-11}\,\mathrm{N\,m^2\,kg^{-2}}$ is Newton's gravitational constant. Dwell on two features of this formula, because they recur, in modified form, for every other force in this chapter.

First, the force is an **inverse-square law**: double the distance and the force drops by a factor of four. This is not an arbitrary choice; it is what one gets whenever something is emitted in all directions from a point source and spreads out over the surface of an ever-growing sphere, whose area grows as r^2 exactly. The same $1/r^2$ structure reappears for the electromagnetic force, for exactly the same geometrical reason, and it is a signature of any force carried by a massless mediator.

Second, $F \propto m_1 m_2$: gravity couples to mass (more precisely, to mass-energy; light itself is deflected by gravity, as in gravitational lensing, Section 1.2.3, even though a photon has no rest mass). There is no such thing as a mass with zero gravitational coupling, which is part of why gravity, among the four forces, is the one every particle without exception must feel.

1.2.2 A worked example: how weak is gravity, really?

Let's actually do the comparison once, rather than simply being told that gravity is weak. Take two protons separated by $1\,\mathrm{fm} = 10^{-15}\,\mathrm{m}$, roughly a nuclear distance, and compare the gravitational force between them with the electrostatic (Coulomb) force between them. Using $m_p = 1.67 \times 10^{-27}\,\mathrm{kg}$ and $e = 1.6 \times 10^{-19}\,\mathrm{C}$,

$$\frac{F_{\mathrm{grav}}}{F_{\mathrm{EM}}} = \frac{G m_p^2 / r^2}{e^2 / 4\pi\varepsilon_0 r^2} = \frac{4\pi\varepsilon_0 G m_p^2}{e^2} \approx 8 \times 10^{-37}. \tag{1.2.2}$$

Notice that the r^2 cancels completely: this ratio does not depend on how far apart the two protons are, only on their intrinsic properties. Gravity is roughly 10^{36} times weaker than electromagnetism between two protons, at any distance. This number, not a vague "gravity is weak", is the actual content of the claim made throughout particle physics that gravity can be ignored at the particle level; it is simply too small by 36 orders of magnitude to matter next to the other three forces, for any two individual particles. Gravity only becomes significant for macroscopic bodies because mass, unlike electric charge, has only one sign: every proton, electron and neutron in the Earth contributes gravitationally with the same sign, so the effect accumulates without cancellation across $\sim 10^{50}$ particles, whereas the

electromagnetic contributions of equal numbers of positive and negative charges cancel almost exactly.

Two defining properties of gravity

- **Always attractive.** Unlike electric charge, mass has no negative counterpart to produce repulsion (there is no such thing as negative mass as far as any experiment has shown), so gravity only ever pulls, never pushes.
- **Infinite range.** The force never drops to exactly zero at any finite distance; it is simply very small far away. This, as noted above, is tied to the photon-like masslessness of its hypothesised carrier, the graviton.

Terms used here: **mass–energy** (Section 1.2.3 below) is the relativistically correct source of gravity, mass m and energy E being interchangeable via $E = mc^2$, so that even energy with no rest mass at all, such as a photon's, still gravitates and is gravitationally deflected. **Range** means the distance beyond which a force's effect becomes negligible; "infinite range" does not mean the force stays strong forever, only that $F \propto 1/r^2$ never actually reaches zero at any finite r, merely becoming smaller and smaller.

Newton's single inverse-square law is remarkably productive: it explains why an apple falls, why a thrown ball follows a parabola, why the planets trace out ellipses around the Sun with the Sun at one focus (Kepler's first law, which Newton's law actually *derives* rather than merely describes), and it lets one predict exactly when a given comet will return, sometimes centuries in advance. It is also, as the next subsection shows, not quite the whole story.

1.2.3 Where Newton's law breaks down, and Einstein's alternative

One specific, small discrepancy eventually forced a rethink. Mercury's orbit is very close to an ellipse, but not exactly: the ellipse itself slowly rotates, so that the point of closest approach to the Sun (the perihelion) shifts position by about 43 arcseconds per century beyond what Newtonian gravity, including the perturbing effect of every other known planet, could account for. This is a tiny effect, but it was measured precisely, and it persisted no matter how carefully the Newtonian calculation was redone.

In 1915 Einstein resolved it with general relativity, which reframes gravity entirely: rather than a force acting instantaneously at a distance between two masses, gravity is the curvature of spacetime caused by the presence of mass or energy. A mass does not reach out and pull on a second mass; instead, it bends the local geometry of space and time, and the second mass, left alone, simply travels along the straightest possible path available to it (a **geodesic**) through that curved geometry. Where Newton saw a force, general relativity sees geometry, and Mercury's perihelion precession falls directly out of the (very slightly) non-Euclidean geometry around the Sun, with no extra assumption required.

Predictions unique to general relativity

- **Perihelion precession**, exactly matching Mercury's observed $43''$/century, with no free parameters left to adjust.
- **Gravitational lensing**: light bends around a massive body, because spacetime itself is curved there, not because light has mass and is pulled on Newtonian lines (a calculation Einstein himself first did incorrectly, by a factor of two, before correcting it).
- **Gravitational redshift**: a photon climbing out of a gravitational well loses energy, and since a photon's energy and frequency are tied together ($E = hf$), this shows up as a shift to lower frequency (redder light).
- **Gravitational time dilation**: a clock deeper in a gravitational well (closer to a massive body) runs slow compared with one farther away; this is not a hypothetical curiosity but an effect that must be corrected for in GPS satellites every day.
- **Black holes**: a region where spacetime is curved so extremely that nothing, not even light, can escape once inside a boundary called the event horizon.

Terms used here:

- **Perihelion**: the point in an orbit closest to the Sun (the corresponding farthest point is the **aphelion**); an **arcsecond** ($''$) is $1/3600$ of a degree, so $43''$/century is an exceedingly small angle, invisible without precise, repeated astronomical measurement over many decades.
- **Gravitational redshift**, more quantitatively: a photon climbing from a radius r near a mass M loses a fractional amount of energy $\Delta E/E \approx GM/(c^2 r)$; since $E = hf$ relates a photon's energy E directly to its frequency f (via Planck's constant h), a loss of energy is exactly a drop in frequency, i.e. a shift toward the red (lower-frequency) end of the spectrum.
- **Event horizon**: for a non-rotating black hole of mass M, the event horizon sits at the **Schwarzschild radius**, $r_s = 2GM/c^2$; at $r \leq r_s$, the escape velocity (the speed needed to escape the gravitational pull entirely) reaches or exceeds c, so nothing, light included, can leave.
- **Geodesic** (used in the main text above): the relativistic generalisation of a "straight line", the shortest, or more precisely the extremal, path between two points through curved spacetime; in flat spacetime (no mass nearby) a geodesic is simply an ordinary straight line, recovering Newtonian motion.

For everyday purposes (building a bridge, launching a satellite into low Earth orbit) Newton's law remains an entirely adequate approximation; general relativity's corrections only become numerically important near very strong gravitational fields or at very high precision. The two theories agree almost everywhere, which is itself informative: it means general relativity had to be built to reduce to Newton's law in the weak-field limit, a consistency check every successor theory in physics must pass.

1.3 The electromagnetic force

If gravity explains the large-scale shape of the universe, the electromagnetic force explains almost everything at human scale and below it. The atom itself exists because a positively charged nucleus attracts negatively charged electrons; chemical bonding, and so the whole of chemistry, is electromagnetic in origin; and so are the optical, electrical, magnetic and mechanical properties that distinguish one material from another (conductors from insulators, magnets from non-magnets, solids from liquids and gases, transparent materials from opaque ones). Pause for a moment on just how much this one force is responsible for: nearly every phenomenon you can see, touch or manipulate directly, with your own hands, in a laboratory or a kitchen, traces back to electromagnetism. Only nuclear physics and gravity's astronomical-scale effects lie outside its reach.

1.3.1 Classical description: Coulomb and Maxwell

Electric charges of like sign repel and of unlike sign attract, following Coulomb's law,

$$F = \frac{q_1 q_2}{4\pi\varepsilon_0 r^2}, \tag{1.3.1}$$

another inverse-square law, for the same geometrical reason as gravity (Section 1.2), but now with a force that can be either attractive or repulsive, because charge, unlike mass, comes with both signs. This single difference, the existence of both positive and negative charge, is why electromagnetism can cancel out over large collections of matter (equal numbers of protons and electrons produce zero net charge) while gravity, with its one-signed source, cannot.

Moving charges generate magnetic fields; changing magnetic fields generate electric fields, and vice versa. This mutual generation is the content of **Maxwell's equations** (1861–62), which unify electricity and magnetism into one classical theory and, as a direct consequence, predict that oscillating electric and magnetic fields can propagate through space as a self-sustaining wave, an **electromagnetic wave**, travelling at a speed fixed entirely by two constants that had previously been measured only in separate electrical and magnetic experiments,

$$c = \frac{1}{\sqrt{\varepsilon_0 \mu_0}}. \tag{1.3.2}$$

When Maxwell calculated this speed from already-known electrical and magnetic constants and found it matched the independently measured speed of light to within experimental error, it was one of the great unifications in physics: light itself was revealed to be an electromagnetic wave, not a separate phenomenon.

1.3.2 The quantum picture: QED

Maxwell's equations are a complete *classical* theory: they correctly describe fields and waves, but they say nothing about what happens at the scale of individual photons and individual charged particles, where quantum effects take over. The modern, quantum account is **quantum electrodynamics (QED)**, developed

chiefly by Feynman, Schwinger and Tomonaga in the late 1940s, and it is, by some distance, the most precisely tested theory in the whole of physics: QED's prediction for the electron's magnetic moment agrees with experiment to better than one part in a trillion.

In QED, two charged particles do not simply "feel" each other's field at a distance, the way a classical picture might suggest; they interact by exchanging a mediating particle, the **photon** (including *virtual* photons, which are never directly observed and which, unlike real photons, are permitted briefly to violate the usual energy–momentum relation, within the bounds set by the Heisenberg uncertainty principle).

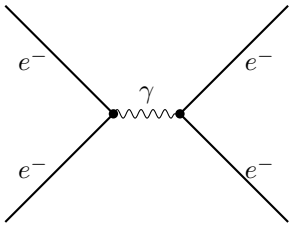


Figure 1.2: Lowest-order QED Feynman diagram for electron–electron (Coulomb) scattering: a virtual photon γ is exchanged between the two electrons, producing the repulsive force between like charges. Time runs upward.

1.3.3 Reading a Feynman diagram: three simple rules

Because Feynman diagrams will recur throughout this course, let's fix, once and for all, exactly how to read one.

1. Time runs in one consistent direction (upward, by convention in these notes); a diagram is read from bottom (initial state, before the interaction) to top (final state, after it).
2. A straight line is a fermion (an electron, a quark, a neutrino); a wavy or coiled line is a boson, the force-carrying particle being exchanged (γ for electromagnetism, gluons for the strong force, $W^{\pm}/Z^0$ for the weak force).
3. Every point where lines meet (a **vertex**) represents one elementary interaction, and must conserve electric charge, and every other relevant conserved quantum number, exactly. A diagram is simply a sequence of such vertices, chained together by the internal (virtual) lines connecting them.

This exchange-particle picture is the general standard-model account of how a force arises at all: wherever two particles attract or repel, the standard model attributes it to the exchange of a specific carrier particle between them, continuously, back and forth, with the net transfer of momentum in each exchange producing exactly the attraction or repulsion a classical force would. For gravity the hypothetical carrier is the graviton (not yet observed, and not successfully incorporated into a full quantum theory, which remains one of the open problems in fundamental physics); for the electromagnetic force it is the photon.

Why the photon's masslessness matters

A rough but genuinely useful rule of thumb, which reappears quantitatively in Section 1.5, is that the range of a force is inversely related to the mass of its carrier particle: a massless carrier (the photon, the graviton) gives an infinite-range force, while a massive carrier confines the force to a short range set by that mass. This single idea is why electromagnetism and gravity reach across the universe while the strong and weak forces act only within, or just beyond, a nucleus.

Where the rule of thumb comes from, in one line: the energy–time uncertainty relation, $\Delta E\,\Delta t \gtrsim \hbar$ (with $\hbar = h/2\pi$ the reduced Planck constant), allows a virtual carrier particle of rest-mass energy Mc^2 to exist only for a time $\Delta t \sim \hbar/(Mc^2)$ before energy conservation must be "paid back"; travelling at up to the speed of light for that time gives a maximum range $R \sim c\,\Delta t \sim \hbar/(Mc)$. Set $M = 0$ (the photon or graviton) and $R \to \infty$; this exact formula is used quantitatively for the pion in Section 1.4 and for the W boson in Section 1.5.

Terms used in this section: a **virtual** particle (as opposed to a **real**, directly detectable one) is a particle appearing only as an internal line in a Feynman diagram, permitted briefly to have an energy and momentum that would not satisfy $E^2 = p^2c^2 + m^2c^4$ for a real particle of that mass, precisely because the uncertainty principle allows such a brief "loan" against energy conservation. A **vertex** is a single point in a Feynman diagram where lines meet, representing one elementary interaction.

1.4 The strong nuclear force

The strong interaction is the strangest of the four forces, in the precise sense that it behaves nothing like the familiar inverse-square forces met so far, so let's build up to it slowly: every one of its odd properties turns out to be necessary to explain something we directly observe.

1.4.1 The puzzle the strong force has to solve

Start from a simple observation: a nucleus with, say, 50 protons packed into a volume only a few femtometres across should be an electrostatic catastrophe. Fifty positive charges crammed that close together repel each other enormously; by Coulomb's law the repulsive force between two adjacent protons at 1 fm is already of order 200 N, a genuinely large force at the scale of a single particle. Nothing in electromagnetism or gravity can counteract this. Yet stable nuclei with dozens of protons plainly exist. Something else, stronger than electromagnetism at this distance and attractive regardless of electric charge, must be holding the nucleus together. That something is the strong force, and (as will emerge) it does its real work one level further down, among the quarks themselves, with the binding of the nucleus as a secondary, residual effect.

1.4.2 Quarks are not optional: what is inside a proton

Protons and neutrons are not elementary: each is a bound state of three quarks. A proton is *uud* (two up quarks, one down); a neutron is *udd* (one up, two down). This is itself remarkable: hit a proton hard enough (in a high-energy collision) and instead of an electron and a nucleus flying apart cleanly, as happens when you ionise an atom, you see evidence of substructure, point-like scattering centres inside the proton, first observed directly in the late 1960s at SLAC in exactly the kind of deep inelastic scattering experiment that confirmed the quark model.

The force that binds the quarks together is mediated by massless **gluons**, and it is governed by a property called **colour charge** (the theory of which is quantum chromodynamics, QCD), with three colours conventionally called red, green and blue (plus their anticolours, antired, antigreen, antiblue). This "colour" is purely a name, borrowed because of a useful analogy to come, and has nothing whatsoever to do with visible colour: it is the QCD analogue of electric charge, except that there are three charges instead of one, each independently able to be positive (a colour) or negative (an anticolour).



Proton (uud) Neutron (udd)

Figure 1.3: Quark content of the proton (uud) and neutron (udd). Each quark carries one of three colour charges (here shown as red, green, blue); the coiled lines stand for the continuous gluon exchange that binds the quarks together.

1.4.3 Why colour must be "colourless" overall

A bound state of quarks is always *colour-neutral* overall: a baryon carries one quark of each colour (red + green + blue), and a meson carries a colour together with its matching anticolour (e.g. red + antired). The analogy with ordinary light is genuinely useful here: mixing red, green and blue light produces white, i.e. no colour at all, and mixing a colour with its complementary colour does the same. QCD requires every *observable*, long-lived particle to be such a "white" combination; any object left with a net, uncancelled colour charge is unstable against further combining, instantly, with other coloured objects until the net colour vanishes.

1.4.4 Confinement: why free quarks are never seen

This colour-neutrality requirement has a dramatic consequence once combined with a second fact about the strong force: unlike the electromagnetic force, whose potential energy between two charges falls off as $1/r$ and eventually becomes

negligible, the strong force's potential energy between two quarks grows roughly *linearly* with separation,

$$V(r) \approx -\frac{4}{3}\frac{\alpha_s}{r} + kr, \tag{1.4.1}$$

for separation r and some constant k (the first term, dominant at short range, looks Coulomb-like; the second, dominant at long range, is new and has no electromagnetic analogue). Pulling two quarks apart therefore costs an ever-increasing amount of energy, rather like stretching a rubber band or a spring, and the field of gluons connecting them concentrates into a narrow, string-like tube rather than spreading out the way an electric field does. Eventually, stretching the "string" further costs more energy than simply creating a new quark–antiquark pair out of the vacuum (recall $E = mc^2$): the string snaps, and the broken ends immediately combine with the newly created quark and antiquark to make two new, colour-neutral hadrons, rather than ever releasing a single free, coloured quark. This is **confinement**: supplying energy to try to pull a quark out of a hadron instead creates a new hadron. No isolated quark has ever been detected, despite extensive, deliberate searches.

Asymptotic freedom: the flip side of confinement

At very short distances, the Coulomb-like term dominates and the effective strength of the strong coupling α_s actually *decreases* as quarks get closer together, the opposite of the electromagnetic coupling's behaviour. This surprising property, **asymptotic freedom** (discovered by Gross, Wilczek and Politzer, awarded the 2004 Nobel Prize), means that quarks behave almost as free, non-interacting particles when probed at very short distances or very high energies, which is exactly the regime deep inelastic scattering experiments access, and exactly why they were able to reveal point-like quarks inside the proton at all.

Terms used here: α_s, the **strong coupling constant**, is the QCD analogue of the fine-structure constant $\alpha \approx 1/137$ that sets the electromagnetic coupling strength (Section 1.3); unlike α, which is very nearly a fixed number, α_s is strongly energy-dependent (physicists say it "runs"), falling from of order 1 at low energy (consistent with the strong force's large relative strength quoted in Section 1.6) to around 0.1 at the energies probed by the LHC. **Deep inelastic scattering** is the technique (firing a high-energy electron at a proton and observing how it scatters) that first revealed quark substructure inside the proton, in the same spirit as, though at much higher energy than, Rutherford's alpha-particle scattering that revealed the atomic nucleus.

1.4.5 The residual strong force, and why the nucleus has a finite size

Colour neutrality between separate hadrons is not perfect at very short range, in the same way that a neutral atom still weakly attracts another neutral atom through van der Waals forces, a residual effect of the electromagnetic force once the main charge cancellation has occurred. The analogous residual effect of the strong force is what binds protons and neutrons into a nucleus: the **residual strong force** (historically called simply the nuclear force), mediated by the exchange of pions (π mesons) between nucleons, acting over roughly 1–3 fm, considerably shorter range than the force acting directly between quarks inside a single nucleon.

This residual force's range is set, by an argument due to Yukawa (1935), directly by the mass of the exchanged pion: a crude but genuinely illuminating estimate uses the energy–time uncertainty relation $\Delta E\,\Delta t \gtrsim \hbar$ to ask how far a virtual pion of mass m_π can travel, at essentially the speed of light, during the brief time $\Delta t \sim \hbar/(m_\pi c^2)$ that energy conservation can be "borrowed" for. The resulting range,

$$R \sim c\,\Delta t \sim \frac{\hbar}{m_\pi c} \approx 1.4\,\mathrm{fm}, \tag{1.4.2}$$

using the pion's mass of about $140\,\mathrm{MeV}/c^2$, matches the observed range of the nuclear force closely, and was in fact how Yukawa *predicted* the pion's existence and approximate mass more than a decade before it was experimentally discovered (1947), a striking example of a force's range and its carrier's mass being two sides of the same coin, a theme already flagged in Section 1.3 and put to further use in Section 1.5.

Because this residual force has a much shorter effective range than the ever-present Coulomb repulsion between protons (which, being carried by the massless photon, never actually switches off, only weakens), there is an upper limit to how large a stable nucleus can be: beyond a certain size, protons on opposite sides of the nucleus are too far apart for the short-range nuclear attraction to overcome their long-range electrostatic repulsion, which is the underlying reason the periodic table terminates in a region of increasingly short-lived, radioactive heavy elements rather than continuing indefinitely.

1.5 The weak nuclear force

The name "weak force" is, if anything, actively misleading, and we should dismantle that misconception immediately rather than carry it through the rest of the course.

1.5.1 Why "weak" is a relative, not an absolute, term

Despite the name, the weak force is far from negligible: within nuclear dimensions it is roughly 10^{25} times stronger than gravity, and it is called "weak" only relative to the strong force, which is itself the odd one out for being so overwhelmingly strong (Section 1.4). In fact, at the very high energies accessible in modern particle colliders, the weak force's intrinsic coupling strength is comparable to the electromagnetic force's; the two unify into a single electroweak interaction at those energies, exactly as Section 1.5.3 below describes. What makes the weak force look weak in everyday, low-energy physics (such as nuclear beta decay) is not a small intrinsic coupling, but something else entirely: the mass of its carrier particle.

1.5.2 Why a massive carrier makes a force look weak

Recall from Section 1.3 that a force's range is tied to its carrier's mass. The same physics, pushed one step further, also controls the force's apparent *strength* at distances shorter than that range. A rough way to see this: the amplitude for exchanging a virtual particle of mass M carrying momentum transfer q includes

a propagator factor $1/(q^2 - M^2c^2)$; at low momentum transfer ($q \ll Mc$, which describes essentially every everyday nuclear process), this factor is suppressed by $1/M^2$ relative to a massless exchange. The $W^{\pm}$ and Z^0 bosons that carry the weak force have masses of about 80 and 91 GeV/c^2 respectively, roughly 80–90 times a proton's mass, which is an enormous mass on the particle-physics scale, and it is this huge mass, not a small coupling, that suppresses the weak force so dramatically at the energies of ordinary nuclear decay.

The Yukawa argument, applied to the weak force

The same estimate used in Section 1.4 for the pion's range applies here too, with the W boson's mass in place of the pion's:

$$R \sim \frac{\hbar}{M_W c} \approx 2 \times 10^{-18}\,\mathrm{m}, \tag{1.5.1}$$

matching the very short range (of order 10^{-17}–10^{-18} m, well inside a proton) at which the weak force is observed to act. Compare this with the pion's much lighter mass (140 MeV/c^2 against 80 000 MeV/c^2), and the roughly 600-fold difference in mass translates directly into a correspondingly shorter range for the weak force relative to the residual strong force.

Terms used here: $\hbar$ is the reduced Planck constant ($\hbar = h/2\pi \approx 1.055 \times 10^{-34}$ J s), the fundamental quantum of action that sets the scale of every quantum-mechanical uncertainty relation, including the energy–time relation $\Delta E\,\Delta t \gtrsim \hbar$ used to derive this formula (Section 1.3). M_W is the mass of the W boson, about 80 GeV/c^2 (equivalently 80 000 MeV/c^2); 1 GeV/$c^2 \approx 1.78 \times 10^{-27}$ kg, so M_W is roughly 86 times a single proton's mass, despite the W boson being, like the electron and quarks, a point-like elementary particle rather than a composite one.

1.5.3 What makes the weak force unique: changing flavour

Its short range and large carrier mass are not, however, the weak force's most physically important feature. That honour belongs to something no other force can do: the weak force is the *only* interaction that can change a quark's (or a lepton's) flavour, for example converting a down quark into an up quark, or vice versa. Neither the strong force (which only ever rearranges existing quarks, never changes a quark's identity) nor the electromagnetic force (which couples to electric charge, leaving flavour alone) can do this. Whenever you see a process in which one type of quark turns into a different type, you are looking at the weak force at work, and nothing else.

1.5.4 Worked example: beta decay, at three levels of detail

The clearest illustration is beta-minus decay, as in the decay of carbon-14, long used for radiocarbon dating:

$$^{14}_{6}\mathrm{C} \;\rightarrow\; ^{14}_{7}\mathrm{N} + e^- + \overline{\nu}_e. \tag{1.5.2}$$

Let's look at this same process at three different levels of resolution, because each level answers a different question.

- **Nuclear level**: a carbon-14 nucleus (6 protons, 8 neutrons) becomes a nitrogen-14 nucleus (7 protons, 7 neutrons), emitting an electron and an antineutrino. This is the level at which the decay was first observed and characterised.
- **Nucleon level**: inside the nucleus, one neutron converts to a proton, $n \to p + e^- + \overline{\nu}_e$. This is the level at which the energetics (the continuous electron energy spectrum that originally motivated Pauli's neutrino hypothesis, Chapter 3) are most naturally discussed.
- **Quark level**: the neutron (udd) and proton (uud) differ by exactly one quark, so the underlying process is a single down quark converting to a single up quark, $d \to u + e^- + \overline{\nu}_e$, mediated by the emission of a W^- boson, which itself (being far too massive to exist as a free particle here, and so existing only fleetingly, virtually) decays almost instantly to the electron and antineutrino actually observed leaving the nucleus.

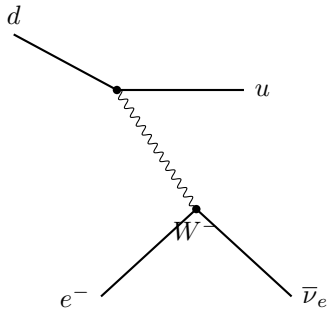


Figure 1.4: Feynman diagram for beta-minus decay: a down quark converts to an up quark by emitting a W^- boson, which itself decays to an electron and an electron antineutrino. Note the single vertex where the quark line changes flavour, the signature of the weak interaction (compare with the QED vertex of Figure 1.1, where the electron line never changes identity).

The mirror process, **beta-plus decay**, converts a proton into a neutron, $p \to n + e^+ + \nu_e$, emitting a positron and a neutrino (rather than an antineutrino), mediated this time by a virtual W^+ boson and an underlying $u \to d$ quark transition; it occurs inside a nucleus whenever energetically allowed, i.e. whenever the parent nucleus is heavy enough relative to the daughter to supply the positron's rest mass and the reaction's kinetic energy.

1.5.5 Electroweak unification

Although the weak and electromagnetic interactions look utterly different at everyday, low energies, one carried by a massless photon with infinite range and the other by extremely massive bosons with a minuscule range, they are, in a precise mathematical sense, two aspects of a single **electroweak** interaction, unified in the electroweak theory of Glashow, Weinberg and Salam (1960s, Nobel Prize 1979). At energies well above the W and Z masses (roughly 100 GeV and beyond, routinely reached at the Large Hadron Collider), the electroweak symmetry is restored and the forces behave as a single, unified interaction with comparable strength; at the low energies of everyday life, that symmetry is "broken" (by the Higgs mechanism,

which also gives the W and Z bosons their mass in the first place, while leaving the photon exactly massless), and the single electroweak force splits into the two very differently-behaving forces we actually observe.

1.5.6 Why stars need the weak force

The weak force is not some exotic add-on to stellar physics; it is essential to it. In the proton–proton chain that powers the Sun, two protons fusing to form a deuterium nucleus (one proton, one neutron) necessarily requires one of the two protons to convert into a neutron, which is only possible via the weak interaction, $p \to n + e^+ + \nu_e$. Without the weak force, this very first step of stellar fusion could not occur at all, hydrogen could not fuse into deuterium, and the entire chain of reactions that builds up helium, and eventually (in heavier stars) all the way to iron and beyond, would simply never get started. Every heavier element in the periodic table, including every atom in your body heavier than hydrogen, ultimately owes its existence to this one weak-force step occurring, against formidable odds, inside a star.

1.5.7 Parity violation: a preview

A further, genuinely startling distinguishing feature of the weak force, examined in full experimental detail in Chapter 4, is that it is the only one of the four forces that violates **parity conservation**: the strong and electromagnetic interactions give identical results under mirror reflection (a process and its exact mirror image are both equally allowed), but the weak interaction does not, revealing a built-in left–right asymmetry woven into the fabric of nature at a fundamental level. This single fact, confirmed decisively by the Wu experiment in 1956, was one of the great shocks of twentieth-century physics, precisely because every other force, and every piece of everyday intuition, had suggested that left and right simply could not be told apart by any law of nature.

1.6 Summary and comparison

With all four forces now built up individually, let's step back and set them side by side: the comparison is more informative than any single force examined alone: almost every one of the differences below is a direct consequence of a single underlying property, the mass (or masslessness) of each force's carrier particle, already flagged repeatedly through Sections 1.2–1.5.

1.6.1 The comparison table

Relative strengths are the conventional order-of-magnitude figures quoted at the scale of a nucleus, where all four can be compared on the same footing (Section 1.2.2 worked through the gravity-versus-electromagnetism entry explicitly; the others follow the same logic).

Force	Mediator	Range	Rel. strength	Acts on
Strong	gluon	$\lesssim 1$ fm (residual: ∼1–3 fm)	1	colour charge (quarks, gluons)
Electromagnetic	photon	infinite	$\sim 10^{-2}$	electric charge
Weak	$W^{\pm}$, Z^0	$\sim 10^{-17}$–10^{-18} m	$\sim 10^{-6}$	flavour (quarks, leptons)
Gravity	graviton (hypothetical)	infinite	$\sim 10^{-38}$	mass–energy

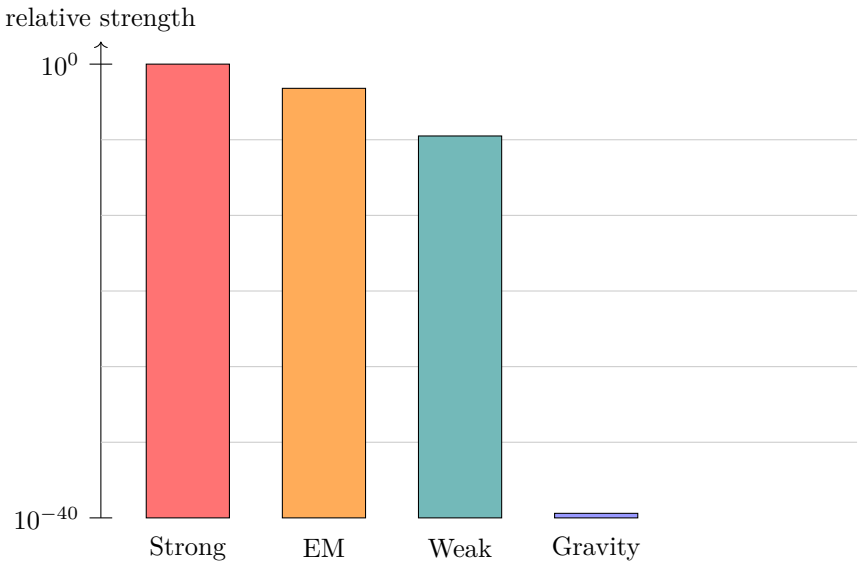


Figure 1.5: The four forces' relative strengths at nuclear distances, on a logarithmic scale (each unit on the vertical axis is a further factor of $\sim 10^{6.3}$). Gravity's bar is not drawn too thin by accident: at this scale it genuinely is, visually, indistinguishable from zero, which is the whole point of the comparison in Section 1.2.2.

1.6.2 A diagnostic habit worth building

One of the most useful skills this chapter can leave you with is the ability to look at an unfamiliar reaction and immediately identify, from its structure alone, which force is responsible, well before any detailed calculation. Three quick questions do most of the work:

1. **Does any quark change flavour?** If yes, the weak force must be involved (Section 1.5); no other force can do this.
2. **If not, is colour charge involved** (are quarks or gluons present, and is the process fast, with a short-lived resonance)? Then it is most likely the strong force (Section 1.4).

3. **If neither**, and electric charge or photons are involved, it is electromagnetic (Section 1.3); if only mass-energy and large, slow, macroscopic bodies are involved, it is gravity (Section 1.2).

This same logic is exactly what let physicists diagnose the strangeness puzzle in Chapter 8 (particles produced quickly, in pairs, by the strong force, but decaying slowly, individually, by the weak force) purely from timescales and production patterns, well before the underlying quark-flavour bookkeeping was understood.

The exchange-particle picture, in one line

In the standard model every force is the result of exchanging a carrier particle: the graviton for gravity (not yet observed), the photon for electromagnetism, gluons for the strong interaction, and the $W^{\pm}/Z^{0}$ bosons for the weak interaction. Only the weak interaction can change a quark's or lepton's flavour, and only the weak interaction violates parity; the strong interaction is unique in having a carrier strong enough to confine its own charge, so that it alone produces a force whose potential energy grows with distance rather than falling off.

A one-line reminder of each term, gathered from Sections 1.2–1.5: **flavour** is which of the (up to six) types of quark, or which of the three generations of lepton, a particle is; **parity** is the mirror-reflection symmetry $\boldsymbol{r} \to -\boldsymbol{r}$ (Chapter 4 covers it in full); "confine its own charge" refers to gluons carrying colour charge themselves (unlike the electrically neutral photon), which is what lets gluons attract *each other* and build the linearly-rising, string-like potential of Section 1.4 rather than the falling-off Coulomb potential every other force shows.

The lecture's closing chapter ("Summary and comparison", from 40:02) recaps points already covered in Sections 1.2–1.5 above rather than introducing new material; the comparison table, bar chart and the diagnostic checklist above are my own synthesis and extension of the figures and properties given earlier in the lecture, cross-checked against the standard values quoted for each force.

2 Antiparticles and Dirac Hole Theory

2.1 Antiparticles: the basic idea

Every matter particle in the universe (electrons, protons, neutrons, neutrinos, quarks, ...) has a corresponding **antiparticle**: a particle with exactly the same mass and spin but with every additively conserved quantum number reversed in sign, most obviously electric charge. The positron e^+, for instance, has exactly the electron's mass ($0.511\,\mathrm{MeV}/c^2$) and spin ($\frac{1}{2}$), but charge $+e$ rather than $-e$.

Let's be clear about what this claim actually asserts, since it is easy to state without appreciating how strong it is. It is not merely that *some* particles happen to come with oppositely charged partners (as, say, the proton and electron happen to have equal and opposite charge, purely by observed coincidence), but that this pairing is a structural, universal necessity, following directly from combining quantum mechanics with special relativity, as Section 2.2 will show for the specific case of the electron. There is no known or conceivable matter particle that lacks an antiparticle partner.

Naming conventions

- Electron e^- $\leftrightarrow$ positron e^+ (a dedicated name, not "antielectron", for historical reasons: it was discovered first, by Carl Anderson in 1932, in cosmic-ray tracks curving the "wrong" way in a magnetic field, well before the general naming convention "anti-X" became standard).
- Proton $p \leftrightarrow$ antiproton $\overline{p}$ (opposite charge, discovered 1955, at the Bevatron, a machine built specifically with enough energy to produce one).
- Neutron $n \leftrightarrow$ antineutron $\overline{n}$: the neutron is electrically neutral, so the antineutron differs not in charge but in baryon number ($B = +1$ versus $B = -1$) and in the (opposite) arrangement of quark content, $udd \to \overline{u}\,\overline{d}\,\overline{d}$. This is an important lesson in itself: "neutral" does not mean "has no antiparticle"; it only means the antiparticle happens to share the charge of 0.
- A neutral particle can, in fact, be its *own* antiparticle (a photon, $\gamma = \overline{\gamma}$, and, as far as is known, the hypothetical graviton) or genuinely distinct from it (a neutron, $n \neq \overline{n}$, as just discussed); which case applies depends on whether the particle carries *any* conserved charge-like quantum number at all, even one with value zero that could in principle still distinguish it from its mirror. A photon carries none; a neutron carries baryon number.

Terms used here: **baryon number** B (defined and used systematically from Chapter 8 onward) is $+1$ for any baryon (a proton, a neutron), -1 for its antiparticle, and 0 for everything else; it is exactly this sign flip, not electric charge (which is already 0 for both n and $\overline{n}$), that distinguishes a neutron from an antineutron. The **Bevatron** was a particle accelerator at Lawrence Berkeley National Laboratory, deliberately designed to reach just enough collision energy to produce an antiproton via pair production (below), confirming a specific, falsifiable theoretical prediction of its minimum required energy.

2.1.1 Annihilation and pair production

When a particle meets its antiparticle they can **annihilate**, converting their entire rest mass and kinetic energy into other particles, most commonly a pair of photons,

$$e^- + e^+ \to \gamma + \gamma, \tag{2.1.1}$$

(two photons rather than one, because a single photon could not simultaneously conserve both energy and momentum in the electron–positron centre-of-mass frame, where the total momentum is exactly zero but the available energy, $2m_e c^2$, certainly is not). Conversely, a sufficiently energetic photon (or other process) can undergo **pair production**, creating a particle–antiparticle pair out of pure energy, provided only that enough energy is available to supply the rest mass of the pair ($E \geq 2mc^2$) and that every conserved quantity, charge, baryon number, lepton number and so on, sums to exactly the same total before and after. This last requirement is automatically satisfied whenever a particle is created *together with* its own antiparticle, since by definition their contributions to every conserved charge cancel: a particle and its antiparticle always sum to zero charge, zero baryon

number, zero lepton number, whatever the specific quantum numbers involved. This is, in a real sense, *why* antiparticles have to exist in a world governed by conservation laws and by $E = mc^2$ together: energy can freely convert into mass, but only in charge-neutral packages, and an antiparticle is nature's way of making a package that is guaranteed to be charge-neutral once paired with its partner.

2.2 Dirac hole theory

Antiparticles were not an experimental afterthought, discovered by accident and only later fitted into theory; they were a theoretical necessity, predicted before they were ever seen. Dirac's relativistic wave equation for the electron (1928) was built to satisfy two requirements simultaneously that had, until then, resisted being combined: consistency with quantum mechanics, and consistency with special relativity. The equation he found does satisfy both, but it comes with a strange, inescapable feature baked directly into its mathematical structure.

2.2.1 The problem: negative energy

Special relativity relates a particle's energy and momentum by

$$E^2 = p^2c^2 + m^2c^4. \tag{2.2.1}$$

Taken at face value, solving for E gives *two* branches for every value of momentum p:

$$E = \pm\sqrt{p^2c^2 + m^2c^4}. \tag{2.2.2}$$

In ordinary, non-relativistic quantum mechanics this square root never arises (energy appears only to the first power in the Schrödinger equation), so the issue never surfaces. But Dirac's equation, being built to be relativistic from the outset, inherits the E^2 relation directly, and so it inherits both branches as genuine, mathematically valid solutions. The negative-energy branch cannot simply be discarded by hand (that would not be a solution of the equation at all, merely a restriction imposed from outside it), and it cannot be ignored as unphysical either, because quantum mechanics insists that a system will, given the chance, transition to *any* available lower-energy state, emitting a photon to carry away the energy difference. If negative-energy states exist and are genuinely empty, every positive-energy electron in the universe should cascade downward through them, one photon at a time, without limit: there would be no stable ground state for any atom, and matter as we know it could not exist. Nature plainly does not behave this way, so something in this reasoning has to give.

2.2.2 The Dirac sea

Dirac's resolution, in 1930, was as bold as the problem demanded: the negative-energy states are not empty at all. They are already, and always have been, *completely filled*, forming an infinite "sea" of negative-energy electrons, one in every available state. The Pauli exclusion principle, which forbids two identical fermions from occupying the same quantum state, then does the rest of the work:

no positive-energy electron can fall into the sea, however favourable the energetics look, simply because every single state down there is already taken. The sea itself, being uniformly and completely filled everywhere and at all times, produces no observable effect of its own: it is, by construction, the reference vacuum, the "nothing" against which everything else is measured.

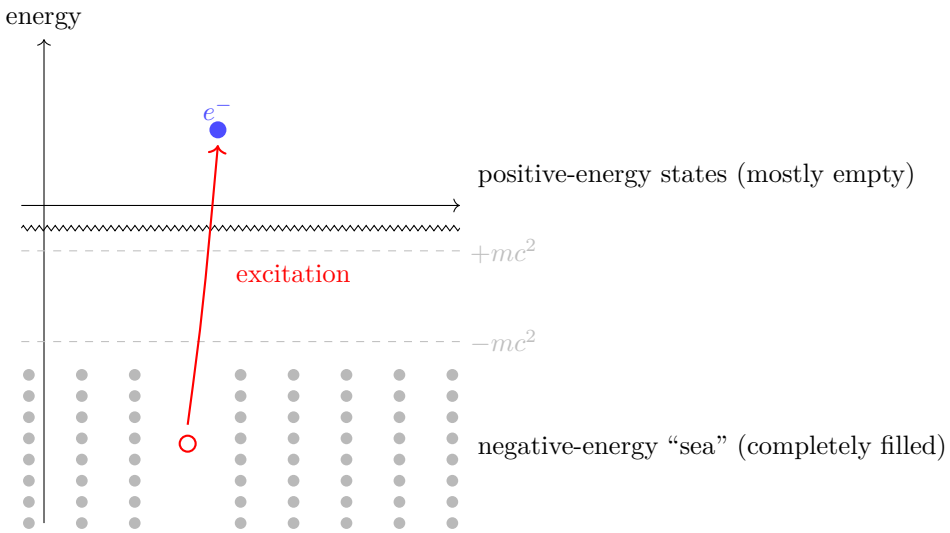


Figure 2.1: The Dirac sea. Every negative-energy state is filled (grey dots); a positive-energy electron (blue) occupies an otherwise nearly-empty positive-energy continuum. Exciting a sea electron upward (red arrow) leaves a hole (open red circle) behind, which behaves as a positive-energy, positively charged particle: the positron.

The hole: an antiparticle

If enough energy, at least $2mc^2$ (enough to bridge the gap between the top of the filled sea and the positive-energy continuum, as the dashed lines in the figure show), is supplied, an electron can be excited out of the sea into a positive-energy state, leaving behind a **hole**. Relative to the uniformly filled sea, which produces no observable effect, a hole is the one place where an electron is *missing* that should be there; it behaves, to every external observer, as a positive-energy, positively charged particle, with exactly the electron's mass, moving as though it were the electron running backwards through spacetime (a slogan usually attributed to Feynman and Stückelberg, developed from a related but distinct viewpoint a decade or so after Dirac's original hole picture). This is the positron: a hole in the Dirac sea.
Terms used here: the **Feynman–Stückelberg interpretation** is the modern restatement, from quantum field theory, that a negative-energy particle travelling forward in time is mathematically identical to a positive-energy antiparticle travelling backward in time; it is a reinterpretation of the equations, not a claim that positrons literally experience time in reverse. The **positive-energy continuum** and the **negative-energy sea** are simply the two branches of solution, $E = +\sqrt{p^2c^2 + m^2c^4}$ and $E = -\sqrt{p^2c^2 + m^2c^4}$, from the Dirac equation's energy–momentum relation introduced at the start of this section; "continuum" signals that, for a free particle, energies within each branch form an unbroken range, with a gap of $2mc^2$ separating the top of the negative branch from the bottom of the positive one.

This single picture explains both directions of the process introduced in Section 2.1. **Pair production**: a photon (or other sufficiently energetic process) excites an electron out of the sea, simultaneously creating an electron and a hole, i.e. exactly the e^- and e^+ pair observed. **Annihilation**: a positive-energy electron loses energy and falls into an existing hole, filling it back in, and the pair's entire rest energy (at minimum $2m_ec^2$) is released as radiation, exactly the inverse of pair production, and exactly consistent with the two-photon final state worked out in Section 2.1.

2.2.3 Why hole theory was eventually superseded

Hole theory is specific to spin-$\frac{1}{2}$ fermions, precisely because it is the Pauli exclusion principle, a fermion-only rule, that keeps the sea stable against collapse. A theory of antiparticles built this way simply has nothing to say about bosons (the photon, for instance, or the pion), for which no exclusion principle exists and so no "sea" can ever be filled in the first place; yet photons and pions certainly have well-defined antiparticle behaviour too (γ is its own antiparticle, as noted in Section 2.1; charged pions have distinct antiparticles, $\pi^+ \leftrightarrow \pi^-$). Later, quantum field theory supplies a cleaner, more general derivation of antiparticles that applies equally to fermions and bosons alike: in that framework, every particle is an excitation of an underlying quantum field, and the theory's mathematical consistency (specifically, the requirement that causality be respected, so that no signal can be sent faster than light) forces every field to admit excitations of *both* signs of charge, with no sea, no infinite filled states and no exclusion principle required at all. Dirac's hole

theory should therefore be understood as a historically pivotal, physically vivid, and genuinely predictive (it came before the positron's discovery) route to the right answer, rather than as the modern, final explanation of why antiparticles exist.

2.3 Crossing symmetry

A further, and genuinely practical, consequence of combining antiparticles with relativistic quantum field theory is **crossing symmetry**: a single Feynman diagram, and the single underlying mathematical expression (the matrix element) that it represents, describes several physically distinct processes at once, obtained simply by moving a particle from the initial state to the final state (or vice versa) and, in doing so, replacing it there by its antiparticle.

2.3.1 The rule, stated precisely

If a process

$$A + B \to C + D \tag{2.3.1}$$

is described by some Feynman diagram, then "crossing" particle C from the final state to the initial state, which requires replacing it by its antiparticle $\overline{C}$ to keep every conserved quantity balanced, gives a *related* process,

$$A + \overline{C} \to B + D, \tag{2.3.2}$$

governed by the same underlying matrix element, now evaluated in a different kinematic region (a different relationship between the particles' energies and momenta). Crucially, this is not merely a bookkeeping convenience, a trick for writing down several related answers at once; it reflects a genuine, deep mathematical fact, that the matrix element for a process, viewed as a function of the particles' four-momenta, is a single analytic function, and the physically distinct processes above are simply that one function evaluated on different parts of its domain.

2.3.2 Worked example

The same fundamental vertex structure (electron, proton, and a photon, exactly the QED vertex already met in Section 1.3) describes all of the following, related only by crossing particles between the initial and final states:

$$e^- + p \to e^- + p \qquad \text{(elastic electron–proton scattering)}, \tag{2.3.3}$$

$$e^- + e^+ \to p + \overline{p} \qquad \text{(annihilation into a proton–antiproton pair)}, \tag{2.3.4}$$

$$\gamma + p \to e^- + e^+ + p \qquad \text{(photon-induced pair production in the field of a proton)}. \tag{2.3.5}$$

Each of these looks, at first glance, like an unrelated physical process, requiring its own separate calculation, its own separate experiment, and quite different intuitions (scattering, annihilation, and pair production are not obviously the

same phenomenon to most people meeting them for the first time). Crossing symmetry says that, underneath this apparent variety, there is just one function doing all the work; which specific process you see depends only on which energy and momentum region you are asking the function about. This is why a particle physicist, looking at a single Feynman diagram, habitually reads it as encoding an entire family of related reactions rather than just the one literally drawn on the page, and it is why measuring one member of a crossing-related family (say, electron–proton scattering, which is comparatively easy to arrange in a laboratory) can give genuine, quantitative information about another member of the same family (say, pair production, which may be much harder to measure directly).

3 Neutrinos

3.1 The neutrino hypothesis

Few particles in physics have as dramatic an origin story as the neutrino: it was proposed not because anyone saw a trace of it, but because without it, one of the best-established laws in physics appeared to be failing.

3.1.1 A crisis in beta decay

Consider beta decay as it was understood in the 1920s: a nucleus emits an electron, converting itself into a slightly different nucleus one place higher in the periodic table,

$$(A, Z) \to (A, Z+1) + e^{-}. \tag{3.1.1}$$

Treat this, as everyone initially did, as a genuine two-body decay (one nucleus in, one nucleus and one electron out). Energy and momentum conservation, applied to a two-body decay at rest, are extremely restrictive: in the parent nucleus's rest frame, the available energy Q (the mass difference between parent and daughter nuclei, converted via $E = mc^2$) is shared between exactly two bodies, the recoiling daughter nucleus and the electron, and simple kinematics (conservation of both energy and momentum, with two unknowns and two equations) fixes the electron's kinetic energy to one single, specific value for a given decay, just as it fixes the kinetic energy of each product in any two-body decay, such as a pion decaying to a muon and a neutrino.

When this was actually measured (chiefly by James Chadwick, 1914, and refined through the 1920s), the result flatly contradicted the prediction: rather than a single energy, the emitted electrons showed a continuous range of energies, smoothly distributed from zero up to some maximum value equal to the Q-value expected for a true two-body decay. Figure 3.1 sketches the shape of this spectrum.

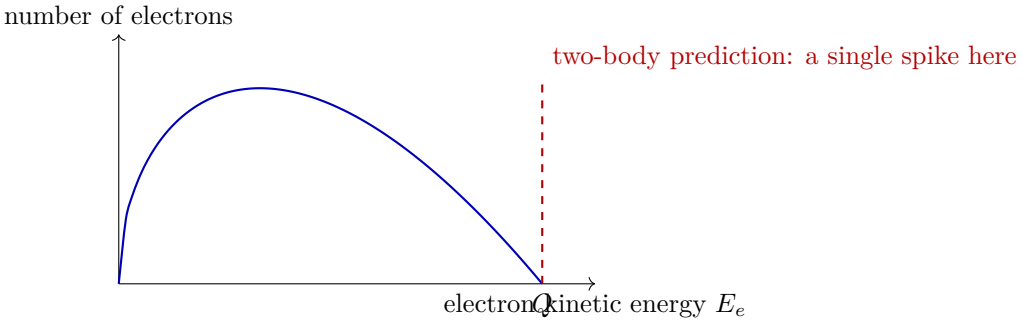


Figure 3.1: The observed beta-decay electron energy spectrum (blue curve): a continuous distribution from 0 up to the maximum value Q, not the single sharp spike at $E_e = Q$ that a true two-body decay would require.

This was not a small, explainable discrepancy; it was a flat contradiction between

theory and experiment, and it was taken extremely seriously. Niels Bohr, no less, was at one point prepared to entertain the possibility that energy conservation itself might only hold *statistically* in nuclear processes, failing in any individual decay while holding on average, a genuinely radical suggestion from one of the founders of quantum mechanics, and a measure of how seriously the crisis was regarded.

3.1.2 Pauli's "desperate remedy"

In December 1930, Wolfgang Pauli proposed a less radical alternative, in a now-famous open letter addressed (with deliberate informality) to "Dear Radioactive Ladies and Gentlemen": rather than abandoning energy conservation, he suggested that beta decay is secretly a *three*-body process, with an additional, extremely light, electrically neutral particle also being emitted, one that interacts so weakly with matter that it had simply escaped every detector built so far. If such a particle shares the available energy Q with the electron, in a continuously variable split between the two of them, a continuous electron spectrum follows immediately and automatically, exactly as observed, with no violation of energy conservation required at all. Pauli himself, in the same letter, called it "a desperate remedy", aware of how bold a proposal it was, and reportedly bet a case of champagne that the particle would never actually be detected (a bet he ultimately lost, gracefully, a quarter of a century later).

3.1.3 Fermi's theory, and getting the name right

Enrico Fermi took Pauli's proposed particle and built the first proper quantitative theory of beta decay around it (1933–34), naming it the **neutrino** (Italian for "little neutral one", deliberately echoing, and distinguishing it from, Chadwick's recently discovered, much heavier neutron). Fermi's theory correctly identified the full, three-body decay as

$$n \to p + e^- + \overline{\nu}_e, \tag{3.1.2}$$

and, importantly, correctly identified the third particle as an electron *antineutrino* $\overline{\nu}_e$ rather than a neutrino ν_e, a distinction that only becomes fully meaningful once lepton-number bookkeeping (worked out in Chapter 8) is in place: assigning $L_e = +1$ to e^- and ν_e, and $L_e = -1$ to e^+ and $\overline{\nu}_e$, the decay $n \to p + e^- + \overline{\nu}_e$ conserves L_e exactly ($0 \to 0 + 1 - 1$), whereas $n \to p + e^- + \nu_e$ would not ($0 \to 0{+}1{+}1$). Getting the particle versus antiparticle assignment right, from theoretical consistency alone, decades before either could be experimentally distinguished from the other, was itself a significant achievement of Fermi's theory.

3.2 Properties of the neutrino

Neutrinos are leptons: elementary, point-like particles (not built of quarks, unlike the hadrons of Chapter 7) that do not feel the strong force at all. Each of its defining properties below rewards understanding rather than just memorising, so let's take them one at a time.

3.2.1 No electric charge

A neutrino carries zero electric charge, so it does not interact electromagnetically: no Coulomb attraction or repulsion, no scattering of light, nothing that couples to the photon (Section 1.3). It interacts only via the weak force (and, in principle, gravity, though at a strength utterly negligible for a single particle, by the argument of Section 1.2.2). This is the direct reason neutrinos pass through ordinary matter in such vast numbers so nearly unimpeded: roughly 60 billion solar neutrinos pass through every square centimetre of your body every second, and the overwhelming majority of them interact with precisely nothing on their way through the entire Earth.

3.2.2 A mass far smaller than any charged lepton's, but (as it turns out) not quite zero

For decades, neutrinos were simply assumed massless in the original formulation of the standard model, both because no experiment had ever measured a nonzero neutrino mass and because the theory could be built in a mathematically simpler way if it assumed exactly zero. The observation of **neutrino oscillations** (Chapter 5), however, showed this assumption to be not quite correct: oscillations can only occur if at least two neutrino mass eigenstates differ in mass, so neutrinos must have some nonzero mass, even though direct measurements still only place an upper bound on its value (of order $1\,\mathrm{eV}/c^2$ or smaller, around half a million times lighter than the electron). Exactly how neutrinos acquire this tiny mass, and why it is so much smaller even than the already-small electron mass, remains an open question in particle physics, one of relatively few places where the standard model is known to need an extension of some kind.

3.2.3 Spin $\frac{1}{2}$, and three flavours

Like the electron, the neutrino has spin $\frac{1}{2}$ (it is a fermion), and there are **three flavours** of it, one associated with each generation of charged lepton (Chapter 6): ν_e, ν_μ, ν_τ. Each flavour is produced and detected alongside its corresponding charged lepton in weak processes: an electron or positron always appears together with ν_e or $\overline{\nu}_e$, a muon or antimuon always with ν_μ or $\overline{\nu}_\mu$, and so on. This pairing is itself a direct consequence of lepton-flavour conservation in the weak interaction (Chapter 8), and it is exactly how experiments identify which neutrino flavour has been produced or detected: not by observing the neutrino directly, which is extremely difficult, but by observing its charged-lepton partner.

3.2.4 A particle defined almost entirely by what it does *not* do

Notice a pattern in this list: every single defining property of the neutrino is either an absence (no charge, no strong interaction, almost no mass) or a bare minimum (just one interaction it does participate in, the weak force). This is precisely why the neutrino was so hard to detect, and why Pauli himself doubted it ever would be: a particle defined by what it *doesn't* interact with is, almost by construction, a particle that leaves almost no trace in any detector built to catch ordinary matter.

How rare is a neutrino interaction, really?

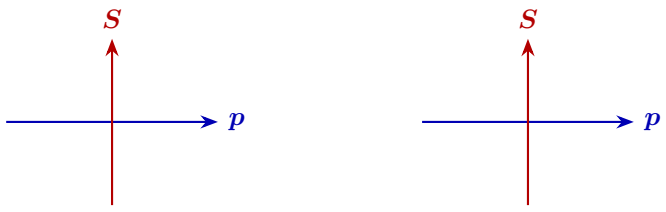
The mean free path of a typical MeV-scale neutrino through solid lead is of order one light-year: that is, a neutrino travelling through solid lead, continuously, would on average need to pass through roughly a light-year's thickness of it before interacting even once. This single fact, more than any formula, conveys just how feeble the weak force's reach is at these energies (consistent with the $\sim 10^{-6}$ relative strength quoted in Section 1.6), and why the first experimental confirmation of the neutrino (Clyde Cowan and Frederick Reines, 1956, using antineutrinos from a nuclear reactor and a detector containing hundreds of litres of liquid scintillator to maximise the tiny interaction probability) came only 26 years after Pauli's original proposal, and required one of the most intense neutrino sources available at the time.

Terms used here: the **mean free path** of a particle through a material is the average distance it travels before interacting once, related to the material's number density n of scattering centres and the interaction **cross-section** σ (an effective target area for the interaction, with units of area) by $\lambda = 1/(n\sigma)$; neutrinos have such a tiny weak-interaction cross-section that λ comes out enormous even through dense solid lead. A **liquid scintillator** is a liquid that emits a flash of light when a charged particle passes through it, the basic detection mechanism Cowan and Reines used to register the rare occasions a neutrino did interact, producing a detectable charged particle in the process.

3.3 Handedness and parity

3.3.1 Helicity: spin relative to motion

Helicity describes the relationship between a particle's spin and its direction of motion. A particle is **left-handed** if its spin points opposite to its momentum (curl the fingers of your left hand in the sense of the particle's spin, and your thumb points backward, against the direction of travel), and **right-handed** if spin and momentum are aligned.



left-handed right-handed

Figure 3.2: Helicity. Left: spin $\boldsymbol{S}$ (red) opposes momentum $\boldsymbol{p}$ (blue): a left-handed particle. Right: spin and momentum point the same way: a right-handed particle. (Both arrows are drawn in a fixed plane purely for clarity; the physical spin and momentum vectors are not generally drawn in 2D, but the sense of alignment or opposition is exactly as shown.)

For a *massive* particle, helicity depends on the observer's reference frame: a

particle moving to the right can always, in principle, be overtaken by an observer moving even faster in the same direction, from whose viewpoint the particle now appears to be moving to the left, while its spin direction (an intrinsic property, unaffected by the observer's own motion) stays fixed. Momentum has reversed, spin has not, so the particle's apparent helicity has flipped, purely because of the observer's choice of frame. For a truly *massless* particle, travelling at exactly c in every frame, no observer can ever catch up to overtake it, so helicity becomes a fixed, frame-independent, genuinely physical property rather than a matter of perspective. Neutrinos, being so nearly massless, behave almost exactly like this massless, frame-independent case in practice, which is what makes the next observation so sharp and clean.

3.3.2 An empirical rule of the weak interaction

Neutrinos come in one handedness only

Insofar as neutrinos are (very nearly) massless, only **left-handed neutrinos** and **right-handed antineutrinos** are ever observed to take part in weak interactions; the mirror-image states, right-handed neutrinos and left-handed antineutrinos, have never been seen to interact weakly, if indeed they exist at all in any detectable sense.

Terms used here, and a touch more structure: for a genuinely massless particle, "left-handed" and "right-handed" are not just two among several possible states, they are the *only* two helicity states a spin-$\frac{1}{2}$ particle can have at all, so the neutrino rule above is really the stronger statement that only one of these two possible states, for the neutrino and its antineutrino each, actually couples to the weak interaction. (For the *massive* charged leptons and quarks that also feel the weak force, both helicity states exist and are observed, but the weak interaction still couples much more strongly to their left-handed components; this handedness-dependent coupling structure is formally called the $\mathbf{V}-\mathbf{A}$ (vector minus axial-vector) structure of the weak interaction, a piece of mathematical machinery this course does not develop in full, but its name is worth having heard before you meet it elsewhere.)

This is a genuinely strange and asymmetric rule to find anywhere in physics: it says that nature, at the level of the weak interaction, has a definite, built-in preference for one specific sense of "handedness" over its exact mirror image. No such preference exists for the strong or electromagnetic interactions, which treat left-handed and right-handed particles completely symmetrically.

3.3.3 The link to parity

This observation connects directly to **parity**, the operation of spatial inversion, $\hat{P} : \boldsymbol{r} \to -\boldsymbol{r}$. Under parity, a particle's momentum $\boldsymbol{p}$ (an ordinary, or **polar**, vector) flips sign, $\boldsymbol{p} \to -\boldsymbol{p}$, exactly as any displacement or velocity would under a mirror reflection. A particle's spin $\boldsymbol{S}$, however, is an **axial** vector (it is built from angular momentum, $\boldsymbol{L} = \boldsymbol{r} \times \boldsymbol{p}$, and the cross product of two polar vectors is axial), and axial vectors do *not* flip sign under parity: $\boldsymbol{S} \to +\boldsymbol{S}$. Putting these

two transformation rules together, parity reverses momentum but not spin, so it necessarily flips a particle's helicity: a left-handed particle becomes right-handed under the mirror, and vice versa.

If the weak interaction respected parity symmetry, left-handed and right-handed neutrinos would therefore have to occur equally often in weak processes, since the mirror-image process (with helicity flipped) would be exactly as likely as the original. They do not: nature's exclusive preference for left-handed neutrinos (and right-handed antineutrinos) is itself direct evidence, just from this one observational fact about neutrino handedness, that the weak interaction cannot respect parity symmetry. This conclusion, reached here almost in passing from neutrino properties alone, is pursued directly, deliberately, and with a completely independent experimental method (not involving neutrino detection at all) in the Wu experiment of Chapter 4.

4 Parity Violation in Beta Decay: the Wu Experiment

4.1 Parity as a symmetry

4.1.1 What a conserved symmetry actually buys you

Before examining parity specifically, let's be clear about why physicists care so much about symmetries at all. A conservation law is not an arbitrary rule handed down from nowhere; by Noether's theorem (1918), every continuous symmetry of a physical system corresponds to a conserved quantity (energy conservation follows from time-translation symmetry, momentum conservation from spatial-translation symmetry, and so on). Parity is a *discrete* symmetry, not a continuous one, so Noether's theorem does not apply to it directly in the same way, but the underlying idea, that a symmetry of the laws of physics implies a corresponding constraint on what can happen, still holds: if the laws of physics are the same whether written in a right-handed or a left-handed coordinate system, then any process and its mirror image must be equally allowed.

4.1.2 Defining parity concretely

The parity operation $\hat{P}$ inverts all three spatial axes simultaneously, $\boldsymbol{r} \to -\boldsymbol{r}$. In practice, this is most easily visualised as the effect of a mirror combined with a 180° rotation (a plain mirror reflection alone only inverts one axis, which turns a right-handed coordinate system into a left-handed one but is not quite the same operation as full parity; however, for the purposes of a real experiment, reflecting an apparatus in a mirror and asking whether the mirror-image process looks physically possible captures the same essential question).

For a physical process, conservation of parity means that the mirror-image process, obtained by reflecting the entire experimental set-up, including every particle's position, momentum, and (axial) spin, according to the correct transformation rules worked out in Section 3.3, is also allowed, and occurs with exactly equal probability: nature, if parity is conserved, does not distinguish "left" from "right" at any fundamental level.

4.1.3 Why this went untested for so long

Until 1956, parity conservation was simply assumed to hold universally, for every interaction, and for good reason: both the strong and the electromagnetic interactions conserve parity to extremely high precision, confirmed by every relevant experiment performed up to that point. Decades of nuclear and atomic physics, built almost entirely on strong and electromagnetic processes, had given no hint of any violation, and "parity is conserved" had settled into the status of an un-

questioned background assumption, rarely examined directly because it had never seemed to need examining.

The question nobody had actually asked

In 1956, Tsung-Dao Lee and Chen-Ning Yang, while investigating an unrelated puzzle in kaon decays (the so-called τ–θ puzzle, in which two particles appeared identical in every measured property except that one decayed to two pions and the other to three, a pattern that seemed to require different parities for the final states, and so, troublingly, for the decaying particle itself), stepped back and asked a more basic question: had parity conservation ever actually been *tested* for the weak interaction specifically, as opposed to simply assumed by analogy with the strong and electromagnetic cases? Searching the existing literature, they found that it had not: every experimental test of parity conservation on record involved only the strong or electromagnetic interaction. They proposed several specific experiments, including the one described in Section 4.2, that would close this gap directly.

Terms used here: a particle's **intrinsic parity** is a further quantum number, $+1$ or -1, describing how its own wavefunction behaves under the parity operation (distinct from, though related to, the orbital parity of a multi-particle final state, which depends on the particles' relative orbital angular momentum too); the τ–θ puzzle's two decay products, two pions versus three, have *different* combined parities by this counting, which is exactly why a single parent particle producing both final states, while conserving parity throughout, looked impossible, unless parity conservation itself was the assumption actually at fault.

Pause on this, because it is a lesson in how physics sometimes advances: not by a new experimental anomaly forcing a rethink (as with the beta-decay spectrum in Chapter 3), but by someone noticing that a long-unquestioned assumption had simply never been checked where it mattered, and then designing the experiment that finally would.

4.2 The Wu experiment

4.2.1 Setting up the test

Chien-Shiung Wu and collaborators (1956, National Bureau of Standards) tested Lee and Yang's proposal using the beta decay of cobalt-60,

$$^{60}_{27}\mathrm{Co} \rightarrow {}^{60}_{28}\mathrm{Ni}^{*} + e^{-} + \overline{\nu}_e, \tag{4.2.1}$$

chosen because cobalt-60 nuclei have a large nuclear spin that can be aligned (polarised) along a chosen axis using a strong external magnetic field. The technical difficulty is that thermal motion of the nuclei, at any ordinary temperature, randomises that alignment almost as fast as the field imposes it; Wu's experiment therefore had to be performed at temperatures of only a few millikelvin, cold enough that thermal agitation no longer disrupts the magnetic alignment, a genuinely demanding cryogenic requirement for 1956.

4.2.2 The logic of the test, spelled out in full

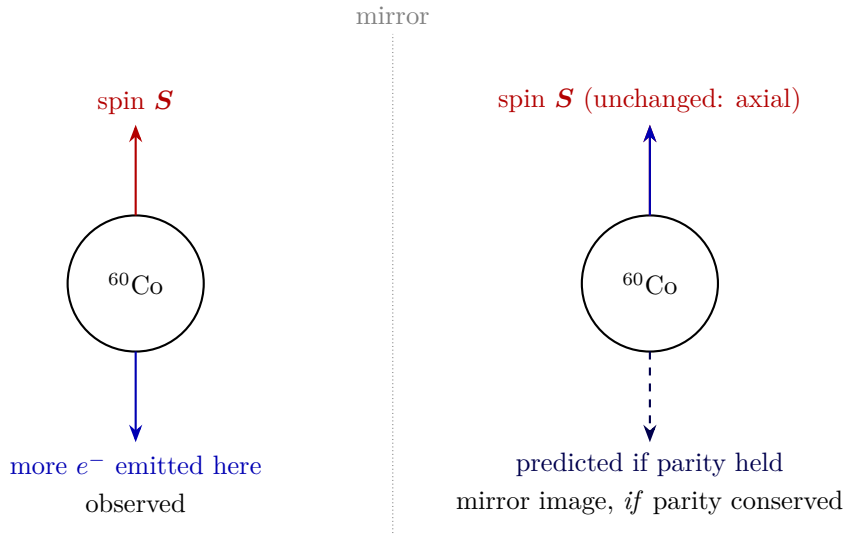


Figure 4.1: The logic of the Wu experiment. Left: the actual result, electrons preferentially emitted opposite to the nuclear spin direction. Right: reflecting the apparatus in a mirror leaves the (axial) spin direction unchanged, so if parity were conserved the mirrored set-up would have to show the same asymmetry, electrons still preferentially opposite the spin; parity conservation makes no prediction that distinguishes "preferentially along spin" from "preferentially opposite spin" once mirrored, since both would have to be equally likely a priori. The experiment distinguishes between these only by comparing the actual emission pattern to what an unbiased, parity-respecting theory would require: equal numbers of electrons along and against the spin.

If parity were conserved, the number of electrons emitted *along* the direction of nuclear spin should equal the number emitted *against* it, with no preference either way. The reasoning: reflecting the entire apparatus in a mirror perpendicular to the spin axis reverses the electron's direction of motion (a true, polar vector, so it flips under reflection) but does *not* reverse the nuclear spin's apparent sense (spin being an axial vector, by the same transformation rule established in Section 3.3). So in the mirror image, electrons that were going *with* the spin now appear to go *against* it, and vice versa; if parity conservation holds, this mirrored situation must be exactly as physically likely as the original, unmirrored one, which is only possible if "with the spin" and "against the spin" were equally likely to begin with. Any observed preference for one direction over the other is therefore, immediately and unambiguously, evidence against parity conservation, requiring no further theoretical input.

4.2.3 The result

Wu's measurement found a clear, substantial asymmetry: electrons were preferentially emitted *opposite* to the direction of the cobalt-60 nuclear spin. Mirroring the experimental set-up manifestly does *not* reproduce an equally likely physical outcome (the mirrored prediction, equal numbers of electrons both ways, flatly

contradicts what is actually observed), so the weak interaction, uniquely among the four fundamental forces, violates parity.

4.2.4 Connecting the asymmetry to Chapter 3's handedness rule

The asymmetry observed here is, in fact, the direct macroscopic signature of the microscopic handedness rule already met in Section 3.3: because the emitted electron antineutrino $\overline{\nu}_e$ must be right-handed (spin along its momentum), and the total angular momentum of the decaying nuclear system must be conserved, the directions in which the electron and antineutrino can be emitted are constrained relative to the nuclear spin axis. Working through the angular-momentum bookkeeping for cobalt-60's specific spin states shows that this constraint forces the electron preferentially in the direction opposite the nuclear spin, exactly the asymmetry Wu measured. In other words, the Wu experiment's macroscopic result and the purely particle-level handedness rule of Chapter 3 are not two independent pieces of evidence for parity violation that happen to agree; they are the same underlying fact, viewed at two different levels of description, nuclear-decay statistics on one hand, individual-particle spin alignment on the other.

4.2.5 The aftermath

The result was confirmed and announced within weeks of Lee and Yang's theoretical proposal, an unusually fast turnaround for an experiment requiring new cryogenic apparatus to be built specifically for the test. Lee and Yang received the 1957 Nobel Prize in Physics for the theoretical proposal, awarded remarkably quickly, within the same year as the discovery itself. Wu, despite performing the decisive experiment without which the theoretical proposal would have remained untested speculation, was not included in that prize; this omission is widely discussed today as one of the more conspicuous examples of an experimentalist's contribution being under-recognised relative to the theorists who proposed the test, and it remains a point regularly raised in discussions of the Nobel Prize's history.

5 Neutrino Oscillations and the Solar Neutrino Problem

5.1 The solar neutrino problem

5.1.1 Why the Sun should produce a precisely predictable flux of neutrinos

The Sun shines because of thermonuclear fusion, principally the proton–proton chain already introduced in Section 1.5: a sequence of fusion reactions, beginning with two protons fusing to deuterium (a step that requires the weak interaction, since it converts one proton to a neutron), and ending, after several further steps, with the production of helium-4. The very first step of this chain,

$$p + p \to d + e^{+} + \nu_e, \tag{5.1.1}$$

necessarily produces an electron neutrino, and because the Sun's total energy output (its luminosity, precisely measured from Earth) is known, and because each full cycle of the proton–proton chain releases a known amount of energy together with a known number of neutrinos, the *total rate* of neutrino production inside the Sun can be calculated from the Sun's luminosity alone, almost independently of the finer details of the solar model. This is a remarkably direct and essentially model-independent prediction: measure the sunlight reaching Earth, and you already know, to good approximation, how many neutrinos the Sun must be producing to supply that much energy.

5.1.2 The experiments, and a persistent shortfall

Starting with Raymond Davis's Homestake experiment (a large tank of chlorine-based fluid located deep underground, in a South Dakota gold mine, specifically to shield it from cosmic-ray backgrounds that would otherwise swamp the tiny neutrino signal) in the late 1960s, and confirmed over the following decades by several later experiments using different detection techniques and different target materials (Kamiokande, using water and the Cherenkov light produced by neutrino-scattered electrons; SAGE and GALLEX, using gallium, sensitive to lower-energy neutrinos than chlorine could detect), the *measured* flux of electron neutrinos from the Sun was consistently found to be only around one-third to one-half of the predicted value.

This was a genuinely puzzling, and genuinely persistent, discrepancy: it showed up across multiple independent experiments, using entirely different detection chemistry and methods, over a period of more than two decades, which ruled out a simple instrumental error specific to any one experiment. It became known as the **solar neutrino problem**, and it sat, essentially unresolved, for roughly thirty years.

Two broad explanations, and why they are hard to distinguish

Faced with a measured flux smaller than predicted, there were, in principle, two very different kinds of explanation: either the standard solar model's prediction of the neutrino flux was simply wrong (perhaps the Sun's core temperature, or its detailed internal structure, differed from the model's assumptions in a way that reduced neutrino production), or the prediction was correct and something was happening to the neutrinos *after* they were produced, somewhere between the Sun's core and the Earth-bound detector, converting some of them into a form the detector could not see. Distinguishing these two possibilities experimentally is not trivial: both predict the same observable symptom (fewer electron neutrinos detected than the solar model predicts), so resolving the puzzle required either an independently cross-checked, more precise solar model, or a detector capable of measuring something other than just the electron-neutrino flux, as Section 5.2 describes.

Terms used here: the **standard solar model** is a detailed numerical model of the Sun's interior (temperature, pressure, density and composition as functions of depth), built from the known physics of nuclear fusion, radiative and convective energy transport, and gravity, and tuned to reproduce the Sun's observed present-day luminosity, radius and age; its neutrino-flux prediction is only one of several independent outputs it must get right simultaneously, including the Sun's measured surface temperature and its helioseismology (the Sun's observed internal oscillation modes), which is why "just adjust the solar model" turned out to be a far less flexible escape route than it might first sound, and why oscillations (Section 5.2) ultimately proved the more productive explanation.

5.2 Neutrino oscillations

5.2.1 Flavour states versus mass states

The resolution turned out to be the second option of Section 5.1: neutrinos produced as one flavour do not necessarily arrive as that same flavour. This is possible because the three *flavour* eigenstates (ν_e, ν_μ, ν_τ, the states that are actually produced and detected in weak interactions, defined by which charged lepton they appear alongside) are not identical to the three *mass* eigenstates (ν_1, ν_2, ν_3, the states that propagate through space with a definite, fixed energy and momentum relationship, i.e. with a definite mass). Instead, each flavour state is a quantum superposition of the mass states, and conversely:

$$|\nu_e\rangle = \sum_i U_{ei}|\nu_i\rangle, \tag{5.2.1}$$

where U is a unitary mixing matrix (the neutrino analogue of the Cabibbo–Kobayashi–Maskawa matrix that mixes quark flavours in the weak interaction), analogous to how a general quantum state can be written as a superposition of energy eigenstates, except that here the two relevant bases, flavour and mass, are simply not the same basis to begin with.

5.2.2 Why propagation makes the mixture change

A neutrino, once produced, propagates through space as a superposition of mass eigenstates, and each mass eigenstate accumulates quantum-mechanical phase at a slightly different rate, because each has a (very slightly) different mass, and therefore a different energy at the same momentum (or a different momentum at the same energy). For a simplified two-flavour case, with mixing angle θ and mass eigenstates ν_1, ν_2 of masses m_1, m_2,

$$|\nu_e\rangle = \cos\theta\,|\nu_1\rangle + \sin\theta\,|\nu_2\rangle, \qquad |\nu_\mu\rangle = -\sin\theta\,|\nu_1\rangle + \cos\theta\,|\nu_2\rangle, \tag{5.2.2}$$

a neutrino produced at $t = 0$ as $|\nu_e\rangle$ evolves after travelling a distance L (at essentially the speed of light, so $t \approx L/c$) into

$$|\nu(L)\rangle = \cos\theta\, e^{-iE_1 L/\hbar c}|\nu_1\rangle + \sin\theta\, e^{-iE_2 L/\hbar c}|\nu_2\rangle, \tag{5.2.3}$$

and projecting this state back onto the muon-flavour basis gives an oscillation probability

$$P(\nu_e \to \nu_\mu) = \sin^2(2\theta)\,\sin^2\!\left(\frac{\Delta m^2 c^4\, L}{4\hbar c\, E}\right), \tag{5.2.4}$$

where $\Delta m^2 = m_2^2 - m_1^2$ and E is the neutrino's energy. The crucial structural point, visible directly in this formula, is that the oscillation probability is only ever nonzero if $\Delta m^2 \neq 0$: *neutrino oscillation can only occur at all if at least two neutrino mass eigenstates have different, nonzero masses*, directly contradicting the standard model's original simplifying assumption of exactly massless neutrinos (Section 3.2).

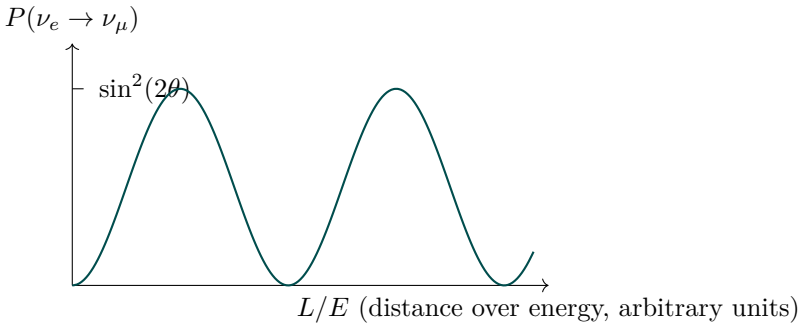


Figure 5.1: The two-flavour oscillation probability as a function of L/E: a neutrino's flavour composition oscillates periodically as it propagates, with an amplitude set by the mixing angle θ and a period set by the squared-mass splitting Δm^2.

5.2.3 How this resolves the solar neutrino problem

Electron neutrinos produced in the Sun's core (Section 5.1) therefore arrive at Earth partly converted into muon and tau neutrinos, in proportions governed by exactly the kind of formula above (generalised to three flavours, and complicated somewhat by matter effects inside the dense solar plasma, the MSW effect, which this course does not pursue in detail but which quantitatively account for the precise observed conversion rate). The original radiochemical solar-neutrino detectors (Homestake,

SAGE, GALLEX) were built to detect only electron neutrinos, via inverse beta decay on a specific target nucleus that only an electron neutrino, not a muon or tau neutrino, has enough of the right coupling to trigger; they were, in effect, blind to the muon- and tau-flavour neutrinos that the Sun's originally-electron neutrinos had partly oscillated into by the time they reached the detector.

Once oscillations are properly accounted for, the *total* flux of all three neutrino flavours, summed together, matches the solar model's original prediction closely, resolving the solar neutrino problem without needing to revise the model of the Sun's interior at all: the Sun was producing exactly as many neutrinos as predicted all along; a large fraction of them had simply changed flavour, and so become invisible, en route. This was confirmed directly and decisively by the Sudbury Neutrino Observatory (SNO, results published 2001–02), whose heavy-water-based detector could, uniquely among the solar-neutrino experiments up to that point, measure the total neutrino flux across all three flavours *as well as* the electron-flavour flux alone, in the same experiment: SNO found the total flux matched the solar model's prediction, while the electron-only flux matched the earlier, lower measurements, directly and quantitatively confirming that oscillation, not a flawed solar model, was the correct explanation all along.

6 Leptons

6.1 Classifying the particle zoo

6.1.1 A history of classification, not just a list

By the mid-twentieth century, particle physics faced something like the situation chemistry faced before Mendeleev: a bewildering number of discovered particles (dozens by the 1950s, hundreds by the 1960s), with no obvious underlying order. Progress, then as now, came from classification before explanation: finding a way to sort the particles into meaningful groups, *before* understanding why those groups existed, often turns out to be the essential first step toward the deeper explanation (exactly as happened with the periodic table itself, which organised the elements by recurring chemical properties decades before anyone understood atomic structure well enough to explain *why* those patterns existed).

The most basic and most useful such split in particle physics is by whether or not a particle feels the **strong interaction** (Section 1.4):

- **Hadrons** respond to the strong force (Chapter 7): they are composite, built from quarks, and this is precisely *why* they feel the strong force, since it is the quarks inside them that carry colour charge.
- **Leptons** do not respond to the strong force at all: they are elementary, point-like particles, not built from anything simpler, and (as Section 6.2 explains) this is precisely because they carry no colour charge to begin with.

This single criterion, feels the strong force or does not, turns out to split the entire particle zoo cleanly into exactly these two camps, with no particle falling ambiguously in between, which is itself a strong hint that colour charge (rather than, say, mass, or some more superficial property) is the deep, structural reason for the split.

6.2 The lepton family

There are six leptons, arranged in three **generations**, each consisting of a charged lepton and its associated neutrino:

Generation	Charged lepton	Neutrino	Charge of charged lepton
1st	electron, e^-	ν_e	$-e$
2nd	muon, μ^-	ν_μ	$-e$
3rd	tau, τ^-	ν_τ	$-e$

together with the corresponding six antileptons ($e^+, \mu^+, \tau^+, \overline{\nu}_e, \overline{\nu}_\mu, \overline{\nu}_\tau$). A useful, and largely accurate, first intuition is that successive generations are essentially

heavier copies of the first: the muon and tau have exactly the electron's charge and exactly the electron's spin, differing from it (and from each other) only in mass, roughly 207 times the electron mass for the muon and about 3477 times for the tau.

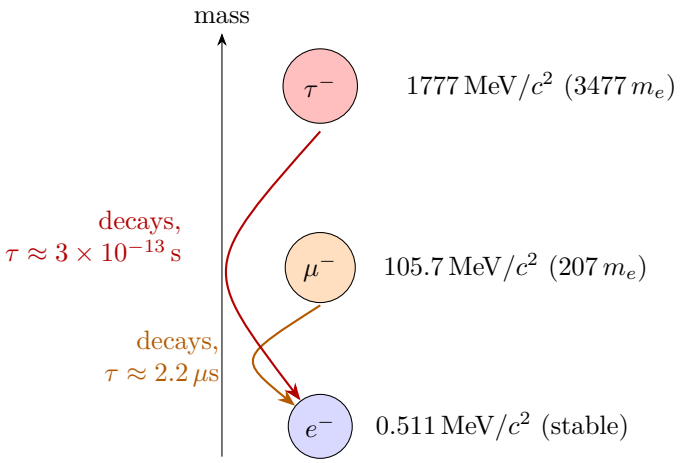


Figure 6.1: The three charged leptons, drawn to show the mass jump between generations (not to scale, but in the right order: each generation is roughly 20 times heavier than the last). Only the electron is stable; the muon and tau both decay, eventually, down to it.

Nobody currently knows why there are exactly three generations, rather than two, or four, or one; this is one of the open questions that the standard model simply takes as an input rather than explaining. It is tempting to treat "three generations" as though it were self-evidently necessary, the way a geometry textbook treats "three dimensions". It is not: as far as anyone can tell, the universe could in principle have had two generations, or five, and the standard model would not object. Nature happened to choose three, and nobody yet knows why.

6.2.1 Stability

Only the lightest charged lepton is stable

The electron is stable: there is no lighter charged particle for it to decay into while conserving both electric charge and lepton number (Chapter 8), so it simply has nowhere to go. The muon and tau are unstable, because each is heavier than the electron and so has energetically available decay channels. The muon decays via the weak interaction,

$$\mu^- \to e^- + \overline{\nu}_e + \nu_\mu, \tag{6.2.1}$$

with a mean lifetime of about $2.2 \times 10^{-6}\,\mathrm{s}$, long enough, at relativistic speeds, for cosmic-ray muons produced high in the atmosphere to reach the Earth's surface in large numbers (itself a classic, direct experimental demonstration of relativistic time dilation). The tau, being heavier still, has several possible decay channels available (including hadronic ones, decaying into pions, since it is heavy enough to do so, unlike the muon), and a much shorter lifetime, around $3 \times 10^{-13}\,\mathrm{s}$: as a rough rule, the more decay channels available to a particle, and the more energy released in each, the shorter its lifetime tends to be, which is exactly the pattern visible across the three generations here.
Terms used here, with the supporting maths: the **mean lifetime** τ of an unstable particle is defined through the exponential decay law, $N(t) = N(0)\,e^{-t/\tau}$, the fraction of an initial population $N(0)$ still undecayed after a time t; it is directly related to the more commonly quoted **half-life** $t_{1/2}$ by $t_{1/2} = \tau \ln 2$. **Relativistic time dilation** is the special-relativistic effect by which a clock (or an unstable particle's own internal "clock") moving at speed v relative to an observer is measured by that observer to tick slowly by a factor $\gamma = 1/\sqrt{1 - v^2/c^2}$; a muon's lifetime in its own rest frame is only $2.2\,\mu\mathrm{s}$, but a cosmic-ray muon moving at, say, $v = 0.999c$ has $\gamma \approx 22$, stretching its lifetime as measured from the ground to roughly $49\,\mu\mathrm{s}$, comfortably long enough to cover the several kilometres from its point of production high in the atmosphere down to a detector at ground level.

6.3 Why leptons are leptons

6.3.1 The real reason, stated precisely

Let's state the defining property of a lepton in its sharpest, most explanatory form, rather than just as a label: a lepton carries no **colour charge** (Section 1.4). Since the strong force couples *exclusively* to colour charge, mediated by gluons that themselves only interact with coloured objects, a lepton simply has nothing for a gluon to couple to; the strong interaction does not fail to act on a lepton by some special exemption or suppression, it does not act at all, for exactly the same reason that the electromagnetic force does not act on a neutral particle: there is no relevant charge present to begin with.

Charged leptons additionally carry electric charge, and so interact electromagnetically (Section 1.3); all six leptons, charged ones and neutrinos alike, interact

weakly (Section 1.5), since the weak force couples to flavour and to a separate "weak charge" that every lepton and quark carries, irrespective of colour or electric charge. This gives a compact, memorable summary: a lepton interacts via gravity, the weak force, and, if charged, the electromagnetic force, but *never* via the strong force, which is exactly the logical complement of the hadron's defining property explored in Chapter 7 (a hadron interacts via the strong force precisely because it is built from colour-charged quarks).

6.3.2 A worked check: can a lepton decay into a hadron?

Let's test this understanding against a concrete question: could a lepton, say a tau, decay directly into a single hadron, conserving charge and energy? The τ^- does in fact have hadronic decay channels, for instance $\tau^- \to \pi^- + \nu_\tau$, which looks at first glance like exactly this. But look more closely at what is actually happening: this is not the tau itself "becoming" a pion; it is the tau decaying via the weak interaction into a (virtual) W^- boson and a ν_τ, with the W^- then converting, again via the weak interaction, into a u and $\overline{d}$ quark pair, which then forms the colour-neutral hadron $\pi^- = d\overline{u}$ (matching the quark content given in Section 7.2 once charge-conjugated appropriately for the specific decay). At no point does the lepton itself acquire colour charge or interact strongly; the weak interaction is doing all the actual work of creating the quarks, and the tau remains, throughout, a particle with zero colour charge that never once couples to a gluon.

6.3.3 Lepton number, previewed

Each lepton generation is separately assigned its own conserved quantum number, the electron, muon and tau **lepton numbers** (L_e, L_μ, L_τ), examined in full in Chapter 8. The muon decay quoted in Section 6.1 conserves each of L_e and L_μ *separately*: the e^- and $\overline{\nu}_e$ produced contribute $+1$ and -1 to L_e respectively, summing to zero, matching the muon's own $L_e = 0$ exactly; the μ^- and ν_μ similarly balance L_μ ($+1 \to +1 + 0$, with the ν_μ carrying the muon's original $L_\mu = +1$ forward). Notice what this means: even though the *total* number of leptons is not conserved in any naive counting sense (one lepton, the muon, became three: an electron and two neutrinos), the generation-specific lepton numbers *are* each separately conserved, which is a far more precise and far more restrictive statement than "lepton number is conserved" would be on its own, and it is this finer-grained conservation law, not a cruder one, that correctly predicts which decays are observed and which (such as $\mu^- \to e^- + \gamma$, discussed in Chapter 8) are not.

7 Hadrons

7.1 Hadrons: particles built from quarks

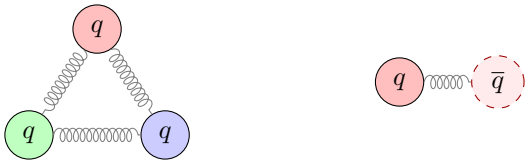
7.1.1 The defining criterion

A **hadron** is any particle that responds to the strong interaction, which, as Chapter 1 and Chapter 6 both already anticipated, means precisely any bound state of quarks held together by the exchange of gluons. This is a satisfying, closed piece of reasoning once all the pieces are in place: a particle feels the strong force if and only if it carries colour charge; a particle carries colour charge if and only if it contains at least one quark; so "hadron" and "contains quarks" are, in fact, exactly the same statement, viewed from two different directions.

7.1.2 Two families, from one simple rule

Hadrons split into two families, according to their quark content, and the split follows directly from the single requirement, already established in Section 1.4, that any observable hadron must be overall colour-neutral:

- **Baryons**: three quarks (qqq), one of each colour, such as the proton (uud) and neutron (udd). Baryon number $B = +1$.
- **Mesons**: a quark and an antiquark ($q\bar{q}$), one colour and its matching anticolour. Baryon number $B = 0$.



baryon: qqq, $B = +1$ meson: $q\bar{q}$, $B = 0$

Figure 7.1: The two ordinary hadron families, in cartoon form: a baryon is three quarks, one of each colour, bound by gluons; a meson is a quark and an antiquark, a colour paired with its own anticolour. Chapter 11 shows these are not the only options QCD allows.

These two families are not the only *conceivable* colour-neutral combinations, merely the two overwhelmingly most common ones. Chapter 11 returns to this point directly, examining tetraquark ($qq\bar{q}\,\bar{q}$) and pentaquark ($qqqq\bar{q}$) states, which satisfy the same colour-neutrality requirement through a different, more exotic combination of constituents.

7.2 Pions: the lightest mesons

The pions are the lightest mesons, built from up and down quarks and antiquarks only:

$$\pi^+ = u\overline{d}, \qquad \pi^- = d\overline{u}, \qquad \pi^0 = \tfrac{1}{\sqrt{2}}\left(u\overline{u} - d\overline{d}\right). \tag{7.2.1}$$

The charged pion's quark content is easy to get backwards, so let's fix it firmly by working the arithmetic through explicitly, rather than just memorising the result: u carries charge $+\frac{2}{3}e$ and $\overline{d}$ carries $+\frac{1}{3}e$ (the antiquark reverses the sign of d's $-\frac{1}{3}e$), so $u\overline{d}$ sums to $+1e$, matching π^+; likewise d carries $-\frac{1}{3}e$ and $\overline{u}$ carries $-\frac{2}{3}e$, so $d\overline{u}$ sums to $-1e$, matching π^-.

The neutral pion's structure, $\pi^0 = \frac{1}{\sqrt{2}}(u\overline{u} - d\overline{d})$, looks stranger at first sight: why a superposition, with a minus sign, rather than simply "$u\overline{u}$" or "$d\overline{d}$" on its own? The short answer is that π^0 is required to be its own antiparticle (it has zero of every charge that would distinguish a particle from its antiparticle, so by the logic of Section 2.1 it can, and empirically does, coincide with its own antiparticle), and neither $u\overline{u}$ alone nor $d\overline{d}$ alone has this property on its own ($\overline{u\overline{u}} = \overline{u}u = u\overline{u}$ happens to work for that term alone, but a single flavour combination does not correctly reproduce the pion's other measured quantum numbers, specifically its isospin assignment within the pion triplet); the specific antisymmetric superposition shown is the one that is simultaneously its own antiparticle *and* carries the correct isospin quantum numbers to sit alongside π^+ and π^- as members of the same triplet (the three pions, together, form an isospin-1 triplet, directly analogous to the spin triplet state in ordinary quantum mechanics).

These are exactly the pions that mediate the residual strong (nuclear) force between protons and neutrons, discussed in Section 1.4, and whose mass, as noted there, Yukawa was able to predict in advance from the observed range of the nuclear force, well before the pion was experimentally discovered.

7.3 Some familiar baryons

Baryon	Quark content	Charge	Note
Proton, p	uud	$+1$	stable (as far as any experiment has shown; see the proton-decay discussion in Section 8.1)
Neutron, n	udd	0	stable when bound in most nuclei, but decays when free, $\tau \approx 880$ s, via $n \to p + e^- + \overline{\nu}_e$ (Section 1.5)
Δ^{++}	uuu	$+2$	an excited state of the nucleon; its existence motivated colour charge directly (Section 1.4 and Section 10.2): three identical u quarks in the same spin state would otherwise violate the Pauli exclusion principle, unless each one carries a different colour
Λ^0	uds	0	the lightest "strange" baryon (Section 8.2), decaying weakly, $\Lambda^0 \to p + \pi^-$

7.3.1 Why the neutron is stable inside a nucleus but not outside it

Pause on an apparent puzzle in the table above: a free neutron decays with a lifetime of under fifteen minutes, yet neutrons inside almost every stable nucleus you have ever encountered (including every atom in your own body heavier than hydrogen) persist for the lifetime of the universe without decaying. The resolution is simply energetics, not any difference in the underlying weak-interaction physics: a free neutron ($939.6\,\mathrm{MeV}/c^2$) is slightly heavier than a free proton plus electron plus antineutrino's rest mass, so the decay $n \to p + e^- + \overline{\nu}_e$ releases energy and is allowed. A neutron bound inside a stable nucleus, however, sits in a lower-energy configuration than the proton it would become if it decayed in place (because converting it would leave the resulting nucleus in a *less* tightly bound, higher-energy configuration overall, once the full nuclear binding energy of both the initial and the hypothetical final nucleus is properly accounted for); with no energetically available final state to decay into, the decay simply cannot proceed, by ordinary energy conservation, exactly the same logic that makes the electron stable (Section 6.1) applied here to a composite nucleus rather than to a single lepton.

7.3.2 From individual particles to a whole spectrum

Hadrons therefore give a natural home to the hundreds of particles discovered through the 1950s and 1960s: each is simply a different allowed combination of a small number of quark flavours, bound by the strong force, differing from one another only in which flavours are present, how they are combined, and in what internal angular-momentum (spin and orbital) state. Making sense of exactly which combinations occur, in what patterns, and why, rather than simply cataloguing them one at a time as this section has done, is the subject of the Eightfold Way (Chapter 9) and the quark model proper (Chapter 10), which together turn this

list of individually-memorised particles into a single, unified, predictive scheme.

8 Lepton Number, Baryon Number and Strangeness

8.1 Lepton number and baryon number

8.1.1 Why bookkeeping rules like this are needed at all

With so many possible combinations of initial and final particles, and so many superficially plausible-looking reactions one could write down on paper, particle physics needs a systematic way to say, quickly and reliably, which reactions actually occur in nature and which do not, well before attempting any detailed dynamical calculation. The conservation laws in this chapter are exactly that: a set of additively conserved quantum numbers, assigned once to every particle, such that any reaction violating one of them can be ruled out immediately, no matter how otherwise plausible it looks (and, just as importantly, any reaction respecting all of them remains a candidate that deserves a real experimental test, even if it looks superficially strange).

8.1.2 Lepton number

Each lepton generation carries its own separately conserved **lepton number**: $L_e = +1$ for e^- and ν_e, $L_e = -1$ for e^+ and $\overline{\nu}_e$, and $L_e = 0$ for everything else, with L_μ and L_τ defined analogously for the muon and tau generations (each particle carries a definite value of all three, usually zero for two of them and ± 1 for at most one).

In every observed interaction, L_e, L_μ and L_τ are each separately conserved, not merely their sum. This is a stronger, more restrictive statement than it might first appear. Let's test it directly on the muon decay already met in Section 6.1 and Section 6.2, $\mu^- \to e^- + \overline{\nu}_e + \nu_\mu$: before the decay, $L_e = 0$ and $L_\mu = +1$ (the μ^- alone); after the decay, $L_e = (+1) + (-1) = 0$ (from e^- and $\overline{\nu}_e$) and $L_\mu = 0 + (+1) = +1$ (from e^-, contributing 0, and ν_μ, contributing $+1$). Both match exactly. This is *why* muon decay must produce both an electron-type and a muon-type (anti)neutrino together, rather than, say, an electron alone: a decay $\mu^- \to e^- + \gamma$ would conserve charge (both sides sum to -1) and conserve energy (for an appropriately energetic photon) perfectly well, but it would violate L_e ($0 \to +1$) and L_μ ($+1 \to 0$) separately, even though their *sum* happens to be conserved in this particular case ($+1 \to +1$). This decay has been searched for with extraordinary sensitivity (it would be a clean signature of physics beyond the standard model if ever seen) and has never been observed, consistent with the separate, generation-specific conservation law rather than the weaker, sum-only version.

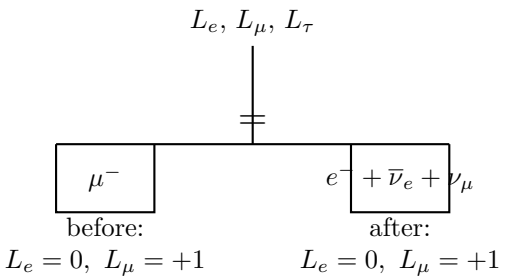


Figure 8.1: The conservation-law check as a balance: whatever total each conserved quantum number has before a reaction, it must have exactly the same total after, pan for pan. Muon decay balances L_e and L_μ *separately*, which is the whole content of Section 8.1's worked check above.

8.1.3 Baryon number

Every baryon (qqq) carries baryon number $B = +1$; every antibaryon ($\overline{q}\,\overline{q}\,\overline{q}$) carries $B = -1$; every meson and every lepton carries $B = 0$ (mesons, being $q\overline{q}$, always have their quark and antiquark baryon-number contributions of $+\frac{1}{3}$ and $-\frac{1}{3}$ cancel exactly). Baryon number is conserved in all observed interactions, to the same standard of "separately and exactly" as lepton number above.

This single rule is what forbids the proton from decaying, even though several energetically allowed decay channels (conserving charge and energy comfortably, such as $p \to e^+ + \pi^0$) can easily be written down: each such channel would reduce the total baryon number from +1 (the proton alone) to 0 (a positron and a pion, both leptons or mesons, both carrying $B = 0$). Proton decay of exactly this kind is, in fact, a generic prediction of grand unified theories, proposals that attempt to unify the strong, weak and electromagnetic forces into a single interaction at extremely high energies, and which typically predict that baryon number is only *approximately* conserved, violated at a very small but nonzero rate. Its absence, despite enormous, dedicated experimental searches (huge underground detectors, built specifically to watch a large number of protons for a very long time, waiting for even a single decay), places a lower bound on the proton's lifetime, if it decays at all, of well over 10^{34} years, vastly longer than the current age of the universe ($\sim 1.4 \times 10^{10}$ years) by more than twenty orders of magnitude. This null result is one of the strongest pieces of experimental evidence constraining, and ruling out, entire classes of otherwise theoretically appealing grand unified theories.

8.2 Strangeness

8.2.1 The puzzle that gave the quantum number its name

Certain hadrons, first seen in cosmic-ray tracks in the late 1940s and studied in detail at accelerators through the 1950s (kaons, the Λ^0 baryon, and others), showed a genuinely puzzling pattern of behaviour, one that took real detective work to untangle. They were always *produced* in pairs ("associated production"), and always on a timescale typical of the strong interaction, extremely fast, consistent with the strong-force cross-sections and reaction rates already familiar from other

strong-interaction processes. Yet each individual particle, once produced, then *decayed* much more slowly, on a timescale typical of the weak interaction, many orders of magnitude slower than the production process that had created it.

This combination, fast production but slow decay, was strange enough, quite literally, that the particles were named for it, and the puzzle itself drove the eventual discovery of a new conserved quantum number.

8.2.2 Gell-Mann and Nishijima's resolution

Murray Gell-Mann and Kazuhiko Nishijima independently resolved the puzzle (mid-1950s) by assigning each quark flavour a **strangeness** quantum number: $S = -1$ for the strange quark s (the historical sign convention here is genuinely awkward, a quirk of how the quantum number was first defined before the quark model itself existed, but it is the universally adopted standard), $S = +1$ for $\overline{s}$, and $S = 0$ for u, d and their antiquarks. A hadron's total strangeness is simply the sum over its constituent quarks: the kaon $K^0 = d\overline{s}$ has $S = 0 + 1 = +1$, and the $\Lambda^0 = uds$ has $S = 0 + 0 - 1 = -1$.

Strangeness is conserved by the strong and electromagnetic interactions, but not by the weak interaction

A strong-interaction process such as

$$\pi^- + p \to K^0 + \Lambda^0 \tag{8.2.1}$$

conserves strangeness overall: the initial state has $S = 0 + 0 = 0$ (neither pion nor proton contains a strange quark), and the final state has $S = (+1)+(-1) = 0$ exactly. Nothing with nonzero strangeness can ever be made *alone* by the strong interaction, since the strong force conserves strangeness and the initial state here has $S = 0$; the K^0 and Λ^0 must be, and are observed to be, produced only together, in exactly the "associated production" pattern that originally prompted the puzzle.

The weak interaction, by contrast, can change quark flavour (Section 1.5), including $s \to u$, and so can change strangeness by one unit, even though the strong and electromagnetic interactions cannot. This is exactly why the Λ^0 (containing a lone s quark among its u, d, s content, with no strange antiquark available to combine with and no strong-interaction channel that conserves strangeness and is energetically open) can *only* decay weakly, e.g. $\Lambda^0 \to p + \pi^-$ (an $s \to u$ transition at the quark level, converting $uds \to uud$ plus a $d\bar{u}$ meson), and hence decays so comparatively slowly: a weak decay, suppressed by the massive W-boson propagator exactly as discussed in Section 1.5, is always far slower than a strong decay would be if one were available, and for the Λ^0, strangeness conservation guarantees that no strong decay channel is available at all.

Terms used here: **associated production** is the pattern, named directly from this puzzle, in which strange particles are only ever created in pairs whose strangeness values sum to the (usually zero) strangeness of the initial state; a **propagator** is the mathematical factor in a Feynman-diagram calculation describing a virtual particle's contribution, roughly $1/(q^2 - M^2c^2)$ for a carrier of mass M and momentum transfer q (already introduced in Section 1.5), whose $1/M^2$ suppression at low q is the precise, quantitative origin of "the weak interaction is slow" whenever a decay is forced to proceed weakly rather than strongly.

8.2.3 A general pattern, not a one-off curiosity

The same pattern recurs for every quark flavour discovered since strangeness: **charm**, **bottomness** and **topness** are each defined completely analogously (a conserved-except-by-the-weak-interaction quantum number tied to the presence of the corresponding quark flavour), and each produces exactly the same diagnostic signature, fast, paired strong-interaction production, followed by slow, individual weak-interaction decay, whenever a particle carries one unit of the corresponding flavour quantum number with no accessible strong-interaction decay channel available. Recognising this pattern is one of the clearest and most reliable diagnostic tools physicists have for identifying, from timing and production data alone, which of the four fundamental interactions is responsible for a given observed process, well before any detailed theoretical calculation is attempted, precisely the same diagnostic habit already recommended at the end of Chapter 1 (Section 1.6.2).

9 The Eightfold Way

9.1 Organising the hadrons

9.1.1 Following the pattern that worked before

By the late 1950s, dozens of hadrons were known, considerably more than the small set of examples listed in Chapter 7, and it became natural to ask the same question chemists had once asked of the elements: might these hadrons too be made of something simpler, combined according to some regular, recurring pattern, in the way atoms are made of nucleons and electrons, and arrange themselves into the periodic table's recurring rows and columns?

Working independently, Murray Gell-Mann and Yuval Ne'eman (1961) found exactly such a pattern: hadrons of the same spin and parity fall into small, regular geometric arrangements when plotted with strangeness S (Section 8.2) on one axis and the third component of isospin I_3 (closely related to electric charge, via $Q = I_3 + \frac{1}{2}(B + S)$, the Gell-Mann–Nishijima formula) on the other. Gell-Mann named the scheme the **Eightfold Way**, after the Buddhist Noble Eightfold Path, because the lightest families of baryons and mesons each contain eight members.

9.1.2 The baryon octet, built up piece by piece

Rather than simply presenting the finished pattern, let's see how it is actually constructed, which makes clear why it has exactly the shape it does. Spin-$\frac{1}{2}$ baryons, built from the three light-ish quark flavours u, d, s, come in the combinations allowed by combining three quarks while respecting the rules already established (Section 1.4, Section 7.1): colour-neutral overall, and, for the particular spin-$\frac{1}{2}$ states that form the octet, a specific allowed symmetry under exchanging any two of the three quarks.

$$n\ (udd) \quad p\ (uud) \qquad\qquad S = 0$$

$$\Sigma^-\ (dds) \quad \Sigma^0, \Lambda^0\ (uds) \quad \Sigma^+\ (uus) \qquad\qquad S = -1$$

$$\Xi^-\ (dss) \quad \Xi^0\ (uss) \qquad\qquad S = -2$$

Figure 9.1: The baryon octet: eight spin-$\frac{1}{2}$ baryons arranged by strangeness S (rows) and charge/isospin I_3 (horizontal position), with quark content shown for each. The proton and neutron sit at the top, $S = 0$, containing no strange quark; each row down adds one strange quark, moving from $S = 0$ to $S = -1$ to $S = -2$.

Read as a grid, the pattern becomes self-explanatory rather than mysterious: moving one row down always means swapping in one more s quark for a u or d (hence S decreasing by one unit each row, matching the strange quark's own $S = -1$ exactly, as defined in Section 8.2), and moving left-to-right within a row

corresponds to the different allowed arrangements of u and d among the remaining, non-strange quark slots (hence changing charge in integer steps, consistent with the Gell-Mann–Nishijima formula). The central position in the middle row is special: it holds *two* distinct particles, Σ^0 and Λ^0, both with quark content uds but differing in how the three quarks' spins are combined internally, a detail the simple (S, I_3) plot alone cannot distinguish, but one that a fuller treatment of the underlying $SU(3)$ symmetry correctly accounts for.

9.1.3 What the pattern was telling everyone

Within this classification, hadrons of a given multiplet differ from one another only in strangeness and I_3, their masses, spins and other properties otherwise closely similar (not identical, since the strange quark is noticeably heavier than u and d, which is why, for instance, the Ξ baryons at the bottom of the octet are measurably heavier than the nucleons at the top, but similar enough that the pattern is unmistakable). This strongly suggested an underlying mathematical symmetry relating them, later identified precisely as an approximate $SU(3)$ **flavour symmetry** (distinct from, though named analogously to, the exact $SU(3)$ *colour* symmetry underlying QCD in Section 1.4 and Chapter 10; flavour $SU(3)$ is only an approximate symmetry of nature, broken by the mass differences between u, d and s, whereas colour $SU(3)$ is believed to be exact), in much the same way that ordinary rotational symmetry relates the different m states of a given angular-momentum multiplet in non-relativistic quantum mechanics, a direct and genuinely useful analogy for anyone already comfortable with that earlier piece of machinery.

9.2 Prediction and discovery of the Ω^-

9.2.1 A second pattern: the decuplet

The decisive triumph of the Eightfold Way came from a second family, the spin-$\frac{3}{2}$ baryons, which form a ten-member pattern (the **decuplet**) rather than an octet, arranged in rows of four, three, two and one as strangeness becomes more negative by one unit per row, exactly the same organising logic as the octet in Section 9.1, just with one more quark-flavour slot available at the bottom because the spin-$\frac{3}{2}$ states permit a symmetric combination that the spin-$\frac{1}{2}$ octet's antisymmetry requirement forbids.

In 1962, nine of its ten members were already known experimentally; only the bottom row's single entry was missing. Gell-Mann noticed the pattern was incomplete at its apex, and, treating the pattern's regularity as more than coincidence, used the near-equal mass spacing observed between the other three rows (each row roughly $145\,\mathrm{MeV}/c^2$ heavier than the one above it, consistent with each successive row swapping in one more, heavier, strange quark) to extrapolate a specific, numerical prediction for the missing tenth particle's mass.

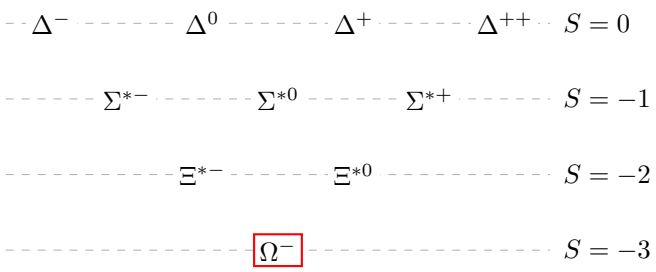


Figure 9.2: The baryon decuplet. Rows of four, three, two and one, strangeness decreasing by one unit per row. In 1962 the bottom row held only a gap where the tenth member should sit; Gell-Mann filled it by prediction, naming the missing particle Ω^- (outlined in red) before it was ever observed.

9.2.2 The prediction, and why it was so bold

He named the predicted particle the Ω^-, with quark content sss (three strange quarks, the only combination left available at the bottom row, $S = -3$), predicted mass around $1680\,\mathrm{MeV}/c^2$, and, from the pattern of decays already seen among the other decuplet members, a predicted decay signature as well: specifically, since no lighter $S = -3$ baryon exists for it to decay into via the strong interaction (which conserves strangeness exactly, Section 8.2), the Ω^- would have to decay weakly, slowly, and with a distinctive sequence of decay products changing strangeness by one unit at a time.

This was a genuinely risky, falsifiable prediction, not a vague gesture at "something should be there": a specific mass, a specific quark content, and a specific, distinctive decay signature, all stated in advance of any experimental confirmation.

9.2.3 The discovery, and what it proved

In 1964 the Ω^- was observed at Brookhaven National Laboratory, in a bubble-chamber photograph showing a particle track with a mass and a decay pattern matching Gell-Mann's prediction closely, essentially exactly as specified in advance. This was comparable, in the history of particle physics, to the periodic table's prediction of undiscovered elements (gallium, germanium, scandium) from gaps in Mendeleev's arrangement, decades earlier in chemistry: in both cases, a classification scheme originally built just to organise already-known entities turned out to have genuine predictive power, correctly specifying the properties of something nobody had yet seen. This is a strong and specific kind of evidence, far stronger than mere after-the-fact accommodation of existing data, and it indicated clearly that the Eightfold Way's patterns reflected real, underlying physical structure rather than a coincidental or merely decorative way of organising already-known particles. It set the stage directly for the quark model proper (Chapter 10), in which that underlying structure is explained concretely, as the various allowed patterns of combining three fundamental, physically real constituents.

10 Quarks: Colour and Flavour

10.1 The quark model

10.1.1 From pattern to substance

In 1964, Gell-Mann and George Zweig independently took the Eightfold Way's success one logical step further. A pattern that merely *organises* known particles is useful, but a pattern that *predicts* unknown ones, as the Eightfold Way had just done with the Ω^- (Section 9.2), practically demands a deeper explanation: why does the pattern have exactly this shape, and not some other? Gell-Mann and Zweig proposed that every hadron is built from a small number of truly fundamental constituents, which Gell-Mann called **quarks** (a name he took, by his own account, from a line in James Joyce's *Finnegans Wake*, "Three quarks for Muster Mark", a nice fit given that the baryon octet and decuplet are built from exactly three quarks each).

10.1.2 Six flavours, discovered one at a time

Six quark flavours are now known, discovered progressively between 1964 (when u, d, s were proposed, though not yet confirmed) and 1995 (when the top quark was finally observed, the last to be found):

Flavour	Symbol	Charge	Approx. mass
up	u	$+\frac{2}{3}e$	a few MeV/c^2
down	d	$-\frac{1}{3}e$	a few MeV/c^2
strange	s	$-\frac{1}{3}e$	$\sim 100\,\mathrm{MeV}/c^2$
charm	c	$+\frac{2}{3}e$	$\sim 1.3\,\mathrm{GeV}/c^2$
bottom	b	$-\frac{1}{3}e$	$\sim 4.2\,\mathrm{GeV}/c^2$
top	t	$+\frac{2}{3}e$	$\sim 173\,\mathrm{GeV}/c^2$

Notice two patterns in this table, rather than treating the numbers as isolated facts to memorise. First, the six flavours come in three pairs, each pair sharing the same charges ($+\frac{2}{3}e$ and $-\frac{1}{3}e$) but increasing sharply in mass: (u, d), then (c, s), then (t, b), exactly mirroring the three generations of leptons already met in Chapter 6 (and indeed the standard model groups a quark pair and a lepton pair together into each of its three "generations", a structural parallel that is almost certainly significant, even though, as noted in Section 6.1, nobody yet knows why there are exactly three). Second, the mass range spans an enormous factor, from a few MeV/c^2 for the lightest quarks to $173\,\mathrm{GeV}/c^2$ for the top, a span of nearly five orders of magnitude, itself an unexplained feature of the standard model (closely related to the unexplained lepton mass hierarchy of Chapter 6) rather than a prediction of it.

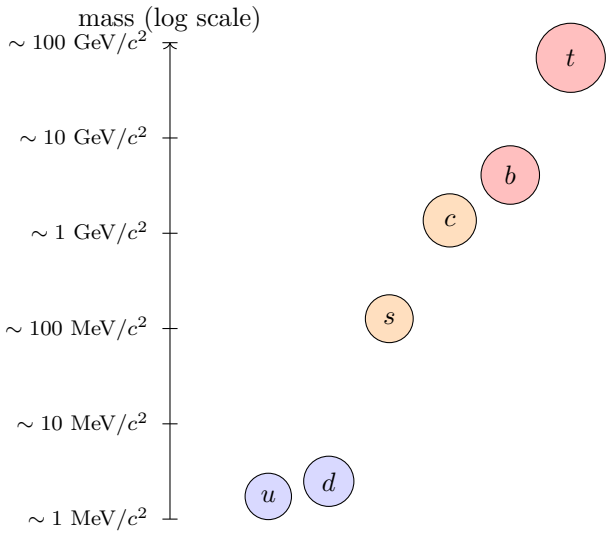


Figure 10.1: The six quark flavours, positioned by mass on a logarithmic scale (circle size also scales with mass, for a second visual cue). Colours group the three generations: (u, d) in blue, (c, s) in orange, (t, b) in red, directly mirroring the three lepton generations of Chapter 6. Notice the top quark sits nearly five orders of magnitude above the up quark.

10.1.3 Why only u and d matter for ordinary matter

The up and down quarks, being by far the lightest, make up ordinary matter (protons and neutrons, and so every atomic nucleus you will ever directly encounter); the heavier flavours appear only transiently, created briefly in cosmic-ray interactions or in particle-accelerator collisions, and then decay rapidly (via the weak interaction, Section 1.5, which alone can change flavour) down to combinations of up and down quarks, the only combinations light enough to be stable or long-lived. This is directly analogous to why, among the six leptons, only the lightest, the electron, is stable (Section 6.1): in both cases, the heavier generations are kinematically unstable, forced by energy conservation to decay down toward the lightest option available.

10.1.4 A genuinely predictive accounting scheme

Every known hadron's quantum numbers, electric charge, baryon number, strangeness, charm, and so on, are reproduced exactly as the simple sum of its constituent quarks' values, a claim already verified directly and repeatedly in earlier chapters: the proton's charge of $+1$ from uud (Section 1.4), the kaon's strangeness of $+1$ from $d\bar{s}$ (Section 8.2), the Ω^-'s strangeness of -3 from sss (Section 9.2). This consistency, across every hadron ever catalogued, is what makes the quark model more than just a relabelling of the Eightfold Way's geometric patterns: it is a genuine, predictive, and now extensively verified accounting scheme for the entire hadron spectrum, reducing hundreds of seemingly independent particles to a handful of combinatorial rules applied to just six underlying building blocks.

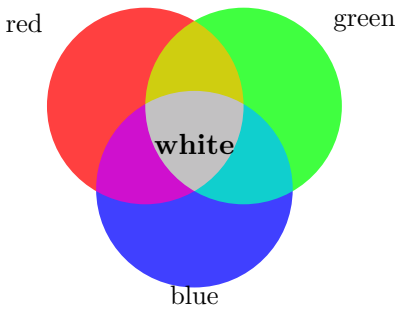
10.2 Why colour was needed

10.2.1 A problem the quark model had from day one

The quark model as first proposed had an immediate, embarrassing problem, already flagged in Section 1.4 and in the baryon table of Section 7.3: the Δ^{++} baryon (uuu, spin $\frac{3}{2}$, the top-right corner of the decuplet in Figure 9.2) requires three identical up quarks, in its ground state, all with the same spin projection (all three quark spins aligned, to produce the baryon's overall spin $\frac{3}{2}$). Up quarks are fermions, and fermions obeying the Pauli exclusion principle cannot occupy an identical quantum state. With nothing else to distinguish the three u quarks, flavour, spin and spatial state all identical, the Δ^{++} should not be able to exist at all, by the exclusion principle alone. Yet it is routinely produced in particle collisions and has been studied in great detail. Something in the quark model, as stated so far, is missing.

10.2.2 The resolution: a new quantum number

The resolution, developed mainly by Oscar Greenberg, Moo-Young Han and Yoichiro Nambu in the mid-1960s, was to give every quark an additional, entirely new quantum number that nobody had previously suspected: **colour charge**, with three distinct values, conventionally called red, green and blue. The three up quarks inside Δ^{++} can then each carry a *different* colour, one red, one green, one blue, making the overall three-quark state properly antisymmetric under exchanging any two quarks (once colour is correctly included in the bookkeeping, alongside flavour, spin and spatial wavefunction), exactly as the Pauli principle requires for identical fermions, while the three quarks remain identical in every other respect: same flavour (u), same charge ($+\frac{2}{3}e$), same spin projection.



red + green + blue = colourless

Figure 10.2: The colour analogy (not a statement about real light, but a genuinely useful mnemonic): just as red, green and blue light combine to give white, a baryon's one quark of each colour combines to give an overall colour-neutral ("white") state. A meson achieves the same colour-neutrality differently, by combining a colour with its matching anticolour.

Colour charge is simultaneously the source of the strong force itself (Section 1.4): quarks interact by exchanging gluons precisely *because* they carry colour, in exact structural analogy to how electric charge sources the electromagnetic force

(Section 1.3) by allowing charged particles to exchange photons. This parallel runs deep: just as there is no electromagnetic interaction between two particles unless at least one carries electric charge, there is no strong interaction between two particles unless both carry colour charge, which, as already noted in Section 6.2, is exactly why leptons, carrying no colour, are entirely invisible to the strong force.

10.2.3 Confinement, revisited with colour explicit

Observable hadrons are always colour-neutral overall, exactly the requirement already used, informally, throughout Chapters 7 and 9: baryons combine one quark of each colour (colourless by the same red+green+blue logic shown in Figure 10.1), and mesons combine a colour with its matching anticolour (equally colourless, by the same logic applied to a colour and its complement rather than to three primaries). No coloured object, no bare quark and no bare gluon, has ever been observed in isolation, which is the precise statement of **confinement**, already introduced and explained dynamically (via the linearly rising potential energy and "string-breaking" picture) in Section 1.4. With colour charge now made fully explicit, confinement can be restated slightly more sharply: it is the requirement that any state left over from the strong-force dynamics with net, uncancelled colour charge is fundamentally unstable, and will continue combining with other coloured objects, converting the energy that would otherwise go into separating it further into the creation of new quark–antiquark pairs instead, until a colour-neutral configuration is reached, exactly the quark-level mechanism underlying every hadron catalogued in this course so far.

11 Tetraquarks and Pentaquarks

11.1 Exotic hadrons

11.1.1 Colour neutrality is more permissive than it first looks

Let's return, one final time, to the single rule that has organised the entire second half of this course: any observable hadron must be colour-neutral overall (Section 1.4, Section 10.2). Chapters 7 and 9 presented this rule alongside only its two most familiar solutions, three quarks of one of each colour (a baryon), or a quark and antiquark of matching colour and anticolour (a meson), and it would be easy to come away thinking those are the *only* solutions QCD permits. They are not. Nothing in the colour-neutrality requirement itself restricts a hadron to exactly three quarks or exactly one quark–antiquark pair; those are simply the two smallest, and so historically the first-discovered, solutions. Larger colour-neutral combinations are equally permitted in principle:

- a **tetraquark**: two quarks and two antiquarks, $qq\overline{q}\,\overline{q}$ (colour-neutral by combining, for instance, a colour-anticolour pair from each of two separate quark-antiquark-like sub-groupings);
- a **pentaquark**: four quarks and an antiquark, $qqqq\overline{q}$ (colour-neutral by combining a baryon-like three-quark colour-neutral core with an additional quark–antiquark pair, or by some other allowed colour pairing).

Such states, collectively called **exotic hadrons**, were not a late or unexpected addition to the quark model's possibilities; they were already anticipated theoretically not long after the quark model itself was proposed (Gell-Mann's original 1964 paper mentions the possibility explicitly, almost as an aside).

11.1.2 Why "theoretically allowed" took fifty years to become "experimentally confirmed"

Despite being permitted from the very beginning, exotic hadrons proved extremely difficult to establish experimentally, so it pays to understand exactly why the gap between prediction and confirmation was so large. Candidate signals were repeatedly claimed over the following decades and then retracted, not through carelessness, but because of a genuine and subtle experimental difficulty: an apparent resonance, a bump in a measured mass spectrum that looks like it could be a new particle, can easily be mimicked by the overlapping, interfering effect of two entirely ordinary, already-known hadrons produced close together in energy. Distinguishing a genuine, bound exotic state from such a kinematic coincidence requires very high statistics (many, many collision events, to be confident a bump is not simply statistical noise) and careful, detailed analysis of the full angular and energy distributions involved, not just the presence of a bump in one simple mass

plot. It took the accumulation of enormous datasets, far larger than were available to earlier experiments, before this distinction could be made with real confidence, which is exactly why the first fully convincing exotic-hadron discoveries did not arrive until the LHC era, discussed next in Section 11.2.

11.2 Confirmed discoveries at LHCb

11.2.1 The experiment that finally settled the question

Firm, statistically overwhelming evidence finally arrived from the LHCb experiment at CERN, analysing the very large samples of b-hadron decays collected during the LHC's operation, precisely the kind of enormous dataset flagged as necessary in Section 11.1.

- **Pentaquarks** (P_c states): first reported in 2015 and confirmed, with substantially higher statistics and finer resolution, in 2019, in the decay $\Lambda_b^0 \to J/\psi\, p\, K^-$. These states have quark content $uudc\bar{c}$: an ordinary uud baryon core (exactly the proton's own quark content, Section 1.4) together with a charm–anticharm pair, rather than five light quarks as the name "pentaquark" might naively suggest.
- **Tetraquarks**: several charm-containing tetraquark states have been reported, including, in 2020, a narrow state consistent with a fully charm tetraquark, $cc\bar{c}\,\bar{c}$, containing no light quarks at all, observed in the di-J/ψ (two J/ψ particles) invariant mass spectrum.

11.2.2 Two competing pictures

Two structural pictures are debated for these states, and the evidence does not yet decisively favour one over the other for every state found. This is a genuinely open area of current research, not a settled textbook fact, and it is refreshing to end a course on a question nobody has fully answered yet.

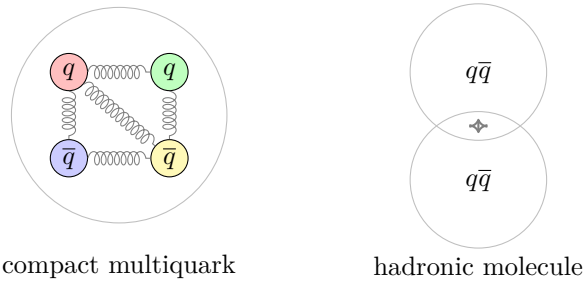


Figure 11.1: Two pictures for an exotic tetraquark state. Left: a compact multiquark, all four (anti)quarks tightly bound together inside one small, colour-neutral object by gluon exchange. Right: a hadronic molecule, two separate, individually colour-neutral ordinary mesons, loosely bound together at a larger separation, directly analogous to how a deuteron is a loosely bound proton–neutron pair rather than a single six-quark object.

A **compact multiquark** has all the quarks tightly bound together by the strong force in a single colour-neutral object, much like an ordinary hadron but with more constituents packed into a similarly small volume; a **hadronic molecule** has two ordinary hadrons (for example a charmed baryon and a charmed meson) weakly bound together at a comparatively larger separation, similarly to how a deuteron is a weakly bound proton–neutron pair held together by the residual strong force (Section 1.4) rather than a single, tightly-bound six-quark object. Both pictures predict states at similar masses in many cases, which is precisely why distinguishing them is difficult: it relies on detailed studies of production rates, decay widths, and quantum numbers (how the state is produced, how quickly and into what it decays, and its spin and parity), an active, ongoing area of analysis rather than a question with a single settled answer applicable to every exotic state discovered so far.

11.2.3 Why this chapter closes the course where it does

As a closing remark for the whole course, notice how directly this final topic connects back to its very first chapter: tetraquarks and pentaquarks exist, and are permitted at all, purely because of the colour-neutrality rule introduced in Section 1.4 to explain something far more mundane, why a proton is uud and not some other combination. The same single rule that explained the simplest hadron in Chapter 1 turns out, fifty years and one major collider later, to also permit, and correctly predict the existence of, some of the most exotic particles ever confirmed. This is a fitting place to end: a reminder that the four forces and their underlying rules, patiently built up one chapter at a time through this course, are not a closed, finished subject but an active, still-developing frontier of research, with LHCb (an experiment this course's author works on directly) continuing to search for further exotic states even now.