

A question-led introduction to particle physics

Matter and its interactions

Mass, charge and the rules of change

David Elliman

Neuro-Symbolic

Modern physics through questions

Volume 2

Preface

Modern physics is fascinating, often surprising, and sometimes hidden behind unfamiliar mathematics. Popular accounts can awaken curiosity without showing enough of the reasoning to satisfy it. University textbooks supply that reasoning, but often assume mathematical fluency that a reader has not yet acquired, or has not used for some time.

At Neuro-Symbolic, we needed material to help able graduates understand the physics relevant to their work. These four books are our attempt to meet that need. They may also interest undergraduates, researchers entering a new field, and readers who want to understand more than a memorable analogy.

The books use questions and answers. We found this approach useful in workshops: a concrete puzzle gives people a reason to examine an idea and discuss their predictions. Each group of questions begins with the physical setting. Explanations use experiments, diagrams and simple arithmetic before introducing symbols. Optional mathematical sections connect that reasoning to the notation used in more advanced texts.

Our aim is to explain what the theories say, why physicists take them seriously, and where their explanations remain incomplete. Established results, assumptions of an example and unsettled proposals are identified separately. A useful calculation need not explain every ingredient that went into it.

We have taken care to check the physics and the examples, but that is no guarantee against error. Mistakes are our responsibility, and we welcome corrections from researchers and readers. The most useful feedback often begins with the exact sentence where an explanation stopped making sense.

Acknowledgement

OpenAI's Codex assisted with organisation, prose, diagrams and teaching checks. Responsibility for the content rests with the author.

How to use this volume

Why does a neutron decay while an electron persists? Why does a neutrino pass through matter so easily? Why do quarks have fractional charges? These questions ask about the rules that distinguish one kind of particle from another and let particles change.

The companion volume *Quantum before symbols* explains amplitudes, states, measurements and spin. Here we use those ideas, but keep the arithmetic visible. An amplitude is a contribution that is combined with other indistinguishable contributions before its squared magnitude gives a probability. A state specifies the quantum preparation. A particle species supplies properties such as charge and spin. A detector record is one outcome of a specified measurement.

At each opening, make a prediction before reading the explanation. The Try questions have answers at the back. A few optional pages introduce notation after the numerical examples. They can be left for another day.

Some questions have firm answers within the Standard Model; others ask why that model has its particular ingredients. The difference matters. Deriving an observable from supplied couplings is a real achievement, but it does not derive the couplings themselves. Rounded numerical examples are labelled as examples, and open questions are left open.

Each part begins by describing the physical setting and the problem that motivates its questions. Read that introduction first. Short transitions then prepare the next question; make the prediction after that context. The optional mathematical pages explain each new operation before using its compact notation.

Contents

Part one: What distinguishes the particles?	6
1 Where does mass come from?	7
2 Why are there three generations?	8
3 Why do quarks have three colours?	10
4 Why do quarks have charges in thirds?	11
5 Why is the electromagnetic coupling close to one over 137?	13
6 Why does a current feel a magnetic force?	14
7 Why do particles bounce instead of merging?	16
8 Which interactions are allowed?	17
9 What is antimatter, and why did matter survive?	19
Part two: What the weak interaction changes	21
10 Why does the weak interaction prefer one handedness?	21
11 Why does a charged weak process change charge by one unit?	23
12 Does beta decay briefly make a lighter W?	24
13 Is a short-lived particle just energy at a frequency?	26
14 Is the weak interaction always weak?	27
Part three: Following neutrinos	29

15 Why are neutrinos hard to detect?	29
16 Why do neutrinos change flavour in flight?	31
17 Why are neutrinos so light, and are they their own antiparticles?	33
18 Where does the energy come from when a neutrino changes chlorine into argon?	34
 Part four: What makes something stable?	 36
19 Does stability require a high barrier?	36
20 Why does the top quark decay rapidly while the proton lasts?	38
21 Why can a bound neutron outlive a free neutron?	39
22 Why do heavy nuclei need extra neutrons?	41
 Part five: Questions behind the familiar answers	 43
23 Why can we not pull one quark out of a proton?	43
24 Can a positron really travel backwards in time?	45
25 What does it mean for the W and Z to eat a field?	46
26 Why these three gauge symmetries?	48
27 Why is strong CP violation so small?	49
28 Why is a light Higgs mass a puzzle?	51
 Optional: Formulas and a colour diagram	 53
29 How does rest energy become a mass formula?	53

30 Why do oscillations reveal differences of squared masses?	55
31 Can three colour arrows explain a proton?	57
For checking: Answers to the Try questions	60
For returning to: Words and further reading	63

Part one

What distinguishes the particles?

Ordinary atoms contain electrons around a nucleus, and the nucleus contains protons and neutrons. Experiments show that protons and neutrons are themselves composite systems of quarks and gluons. A gluon belongs to the strong interaction, which binds quarks. This immediately gives us two kinds of question: which particles exist, and how do their interactions make the structures we observe?

The Standard Model is the field theory that describes the known elementary particles and their electromagnetic, weak and strong interactions. A field has quantities that can vary throughout space. Its independently variable components are called degrees of freedom. In quantum theory, excitations of a field can be prepared and detected as particles in suitable circumstances. Fermions have half-integer spin and obey the exclusion rule discussed in the quantum volume; bosons have integer spin and allow a different pattern of shared occupation. A coupling specifies how strongly particular fields interact. Gravity requires a further part of our account and is discussed in the next volume.

We begin with mass and the particle families, then compare the two kinds of charge called electric charge and colour. Finally we ask how interactions change particles. Some answers follow from well-tested laws. Others, such as why the theory has its observed parameters, remain open. Keeping those two kinds of answer separate is part of understanding the theory.

1 Where does mass come from?

Before reading on. If a proton has the net quark content two up quarks and one down, should adding their individual masses give the proton's mass?

The rest mass of a specified system is determined by its total energy in the frame where its total momentum is zero, when such a frame exists. For a composite system, its rest mass includes the energy of motion, fields and interactions in its rest frame. Adding the masses of isolated constituents can therefore miss most of the answer.

The proton is the important example. Its three valence quarks specify its net quark content: two up quarks and one down after quark-antiquark pairs are accounted for. The full quantum state also contains gluon and quark-pair contributions. They are not three static beads whose masses almost add to a proton. The strongly interacting quark and gluon system accounts for most of the proton's rest energy. Quantum chromodynamics, or QCD, is the theory used to calculate that dynamics.

Elementary fermion masses pose a different question. The electroweak theory describes electromagnetism and the weak interaction together. It includes the Higgs field. A vacuum is the reference state with no added particles in the specified setting; it does not require every field value to vanish. In the usual description, the Higgs field has a nonzero background value even in that state.

Coupling a charged fermion to this background produces a mass term, the part of its equation that fixes its rest energy. The background and the coupling strength both enter the result. The associated change in how the field symmetries appear is called electroweak symmetry breaking. Question 25 explains the field reorganisation more fully. The mechanism permits mass consistently with the theory's gauge structure; it does not predict all those coupling values.

This distinction is visible in the muon. A muon has the same electric charge as an electron but roughly 207 times its mass. The familiar theory accommodates the ratio through different couplings. Merely saying that both interact with the Higgs field does not explain why the ratio is close to that number.

For a useful piece of arithmetic, imagine a bound system whose separated constituents have total rest energy 100 units. If forming the bound state transfers 3 units to the surroundings and leaves the bound system at rest in that frame, its remaining rest energy is 97 units. For example, radiation could carry the energy away with zero net momentum. The mass is correspondingly smaller. This is an illustrative energy account, not a proton model: strongly confined quarks cannot be treated as freely separated building blocks in that fashion.

Mass measurements consequently probe different physics in different systems. For an electron we ask about an elementary mass term. For a nucleus we also need the binding of its constituents. The same energy account connects both questions, while the dynamics that supplies the energy differs.

Try 1. In the illustrative binding account, where did the missing 3 units of rest energy go?

An electron and a muon have the same electric charge but different masses and different decay behaviour. This is one example of a broader repetition in the particle catalogue. We can describe that pattern accurately even while asking why nature contains it.

2 Why are there three generations?

Before reading on. If we arrange known particles into three repeating families, have we explained why there are three?

The observed fermions fit a striking pattern. Each generation contains two quark types, a charged lepton and a neutrino flavour associated with that charged lepton. Leptons are the charged electron, muon and tau, together with their associated neutrinos; they do not carry strong colour charge. The quark entries each have three colour components. Antiparticles also belong in the full account. The following table shows the family pattern without trying to draw all those components.

	First	Second	Third
Quarks	up, down	charm, strange	top, bottom
Charged lepton	electron	muon	tau
Neutrino flavour	electron	muon	tau

The later generations repeat the basic charge pattern while changing the masses. They are not simply excited atomic states of an electron. A muon is an elementary species in the model, and its decay and production are described by that model's interactions.

Much ordinary matter uses first-generation charged fermions. Higher-generation particles are important in energetic collisions and in virtual contributions to processes. Their presence changes which interactions and mixing effects are possible, even if they do not make stable everyday atoms.

The Standard Model describes three generations; it does not provide a generally accepted explanation of why that number occurs. Experiments constrain additional particles very strongly in many cases, but no fourth generation needs a model attached to it. Extra matter with different charges or interactions is not covered by every bound on a straightforward fourth copy.

One historical clue comes from the invisible decay rate of the Z boson, a neutral carrier of the weak interaction. Invisible here means that the decay products, such as neutrinos, escape the detector without its usual visible tracks. It constrains the number of sufficiently light neutrino species with the specified standard coupling. It is not a census of every possible neutral particle at every mass.

Counting can reveal an attractive organisation. To become an explanation, a proposed count must also select the physical representations, permit the observed interactions and account for why competing assignments are absent. Otherwise we have relabelled the catalogue, which can be useful, but have not derived its length.

Try 2. Would a limit on light neutrinos with standard Z couplings automatically exclude a heavy neutral particle with different couplings?

Quarks take part in the strong interaction as well as carrying electric charge. The strong interaction needs its own charge description, called colour. We now ask what the three colour components mean and what follows once that structure is specified.

3 Why do quarks have three colours?

Before reading on. Does a red quark reflect red light, or does colour name a different property?

Colour is a label for the charge structure of the strong interaction. Red, green and blue are convenient names for three components of a quark's colour state. They have nothing to do with the wavelengths that make a painted surface look red or blue.

QCD describes transformations among these components using a symmetry called $SU(3)$. The three-dimensional quark representation is a physical input to that theory, supported by its successful predictions. Here a representation means a specified rule for how the components change under the symmetry. $SU(3)$ is the name of the particular set of allowed colour transformations used by QCD. Once that representation and symmetry are supplied, the associated gluon structure is constrained.

A tempting count starts with three colours and three anticolours, making nine colour–anticolour combinations. They can be reorganised into one colour singlet and an eight-dimensional part. QCD's gluon field belongs to the latter. The arithmetic is $3 \times 3 = 9$ and $9 - 1 = 8$. This is representation counting, not a picture of each gluon as two tiny coloured objects stuck together.

Hadrons are composite particles made from strongly interacting quarks and gluons; protons, neutrons and mesons are examples. Observable isolated hadrons are colour singlets: their full colour state is unchanged by the colour symmetry. A meson can combine quark and antiquark colour into a singlet. A baryon can combine three quarks into a singlet. In a quantum state these are appropriate combinations of colour assignments, not merely a permanent red, green and blue seating plan.

Trying to pull quarks apart does not usually uncover isolated quarks at the ends of a stretched spring. Energy stored in the strong field can produce new quark–antiquark pairs, leading to colour-singlet hadrons. The detected jets in energetic collisions are related to the underlying quarks and gluons through this dynamics.

The question why three? is deeper than the nine-minus-one calculation. That calculation explains the gluon count given $SU(3)$, but does not select $SU(3)$ from all possible gauge theories. An explanation of that choice would have to account for the observed strong interactions as well as the number of labels.

Try 3. Which part of the gluon-count argument is supplied before the arithmetic nine minus one can be used?

The strong colour label does not replace electric charge. Up and down quarks have different electric charges, while familiar isolated protons and neutrons have integer total charges. Adding the constituents' charges is a useful first check on that account.

4 Why do quarks have charges in thirds?

Before reading on. Could two up quarks and one down quark make a particle with charge +1 even though none individually has charge +1?

Measure electric charge in units of the positive elementary charge. An electron has charge -1; an up quark has $+2/3$; a down quark has $-1/3$. The arithmetic of the proton's valence content is then straightforward:

$$\frac{2}{3} + \frac{2}{3} - \frac{1}{3} = 1.$$

For the neutron, one up and two down quarks give

$$\frac{2}{3} - \frac{1}{3} - \frac{1}{3} = 0.$$

The neutron is electrically neutral as a whole, although its internal charged constituents still affect its electromagnetic properties.




Figure 1: Valence charge accounting for a proton and a neutron. The circles are a ledger of charges, not a scale drawing of their internal wavefunctions.

Charge addition is a conservation rule, not an explanation of the fractional inputs. The Standard Model assigns fields particular transformation rules under its electroweak symmetry. One assigned charge is called hypercharge. Electric charge is a particular combination of the electroweak charges, so their assignments must be specified before the electric charges follow. After symmetry breaking, those assignments produce the observed electric charges.

Consistency conditions constrain the assignments. For example, a quantum calculation must not destroy the gauge symmetry that the theory needs for consistency. Such a failure is called a gauge anomaly; the matter contributions must combine so that it cancels. But a claim that charge has been derived still needs to disclose the assumed field content, representation choices and charge normalisation. Those assumptions do real work.

Fractional quark charge is not seen by placing an isolated down quark on a balance. Quarks are confined, and their charge assignments are tested through a network of scattering, production and hadron measurements. The total charge of the familiar isolated hadrons is integer in these units.

A useful record scheme should get the addition right for every allowed process. For example, turning a down quark into an up quark changes the quark's charge by one unit. Something else in the process must carry the compensating change. That requirement will organise beta decay later.

Finding a set of integers or fractions that can name these charges is only the first step. The fields must couple to electromagnetism with those charges, and their interactions must preserve the total. The arithmetic becomes physics through those observable responses.

Try 4. What are the total valence charges of a proton and a neutron in these units?

Knowing the signs of electric charges does not tell us the size of an electromagnetic response. A coupling supplies that strength. The fine-structure constant expresses it as a dimensionless number, so a change of measurement units cannot change the comparison.

5 Why is the electromagnetic coupling close to one over 137?

Before reading on. Would counting exactly 137 objects be enough to predict a measured inverse coupling near 137.036?

The fine-structure constant is a dimensionless measure of electromagnetic coupling. Dimensionless means its numerical value does not change because we switch from metres to centimetres or from one unit of charge to another, provided we transform the whole expression consistently.

At low energy its inverse is approximately 137.036. The nearby integer 137 is memorable, but it is not the measured value. A proposal predicting exactly 137 needs a justified response or correction calculation before it can be compared at high precision.

Here is the scale of that difference. The difference between 137.036 and 137 is 0.036. Compared with 137.036, that is about 0.026 per cent. Small to the eye does not mean small relative to the accuracy of a precision experiment. Numerical resemblance needs an uncertainty and a physical definition attached to it.

The effective electromagnetic coupling also depends on the momentum scale of the process. Quantum fluctuations affect the relation between a parameter in a calculation and a measured response. This scale dependence is called running. It is not a claim that the charge of a stored electron randomly changes from one moment to the next.

A physical determination uses specified observables and a theory connecting them to the coupling. Different experiments provide important cross-checks. A microscopic model must say which quantity it predicts.

A bare parameter is a coefficient before the corrections associated with the chosen short-distance description are combined with it. A renormalised parameter is defined after those corrections have been accounted for, using stated conventions and a scale. An observable is the measured response to which the full calculation must be compared. Mixing those roles can create an impressive but meaningless agreement.

The Standard Model uses the coupling as an input. It predicts many consequences once the input and the other parameters are fixed. Explaining its value from more basic premises remains a further task. The quality of a proposed answer depends on its derivation and tests, not on how many digits can be fitted after the measurement is already known.

Try 5. What is the difference between 137.036 and 137? Why does that difference need comparison with experimental precision?

Electric charge appears in motion as a current. A current-carrying wire can feel a force near a magnet, as in an electric motor. The force depends on the geometry as well as on the amount of charge flowing, which makes it a useful concrete interaction to examine.

6 Why does a current feel a magnetic force?

Before reading on. A straight wire lies parallel to a uniform magnetic field. Must it feel the same magnetic force as an otherwise identical wire placed at right angles?

A current is a flow of charge. Moving charges interact with electric and magnetic fields. For the magnetic part of the force, direction matters: it is perpendicular to both the charge's velocity and the magnetic field. Reverse the charge sign and the force reverses.

In a wire, the charge carriers interact with the material. The resulting momentum transfer can exert a force on the wire as a whole. The field-producing magnet or coil and the electromagnetic field also belong in the momentum account. There is no isolated one-way push with no reaction elsewhere.




Figure 2: A wire segment perpendicular to a field directed into the page. The arrow is conventional current. The force direction is shown for that current and field orientation. Crosses mark the field direction.

For a straight segment in a uniform field at right angles, the familiar wire rule multiplies current, length and field strength. Take a current of 2 amperes, a segment 0.5 metres long and a field of 0.3 tesla. The force magnitude is $2 \times 0.5 \times 0.3 = 0.3$ newtons. These are chosen teaching values, not a measurement of a particular motor.

Rotate the segment until it is parallel to the field. The magnetic force on that straight segment then vanishes in this ideal model. Intermediate angles give intermediate forces. This angular dependence is essential to the operation of motors and other electromagnetic devices.

The magnetic force on an individual point charge is perpendicular to its instantaneous velocity, so that force alone does no work on the charge. This does not mean an electric motor cannot deliver work. A motor includes a circuit, electric fields, mechanical constraints and a power supply. Its energy transfer must be described for the complete device.

Electricity and magnetism are parts of one electromagnetic field. Different inertial observers can divide its effects differently into electric and magnetic components while agreeing on the physical events. The field and the motion of the charges together determine the prediction.

Try 6. What happens to the force in the numerical example if the current is doubled with the geometry unchanged?

Interactions can change motion, bind objects or change the species present. In a collision we prepare incoming states and measure the outgoing ones. Calling the event a bounce may conceal other possible outcomes, so we begin with the conservation checks.

7 Why do particles bounce instead of merging?

Before reading on. If two incoming particles meet, is elastic scattering the only possible outcome?

Particles can scatter, become bound, annihilate with antiparticles or produce new particles when the available energy and interactions permit it. Bounce describes only one class of outcome. Even there, quantum scattering generally predicts a distribution of outgoing momenta rather than one billiard-ball trajectory.

Conservation laws narrow the possibilities. The total charge, energy, momentum and angular momentum must be accounted for. Further selection rules depend on the theory and the states. Passing those checks does not guarantee a large probability: the relevant amplitude may be small or vanish.

For an imaginary collision with 10 units of total energy in its centre-of-momentum frame, a proposed final state requiring 12 units of rest energy is unavailable. A final state requiring 8 units may leave 2 units for kinetic energy, provided momentum and the other rules can be satisfied. This accounting is a threshold check, not an interaction calculation.

Two weak waves in a linear classical medium can pass through each other and add during their overlap. An interacting field theory has additional couplings, which determine whether its excitations scatter or change. This does not abandon the linear evolution of quantum states: the interaction is part of the Hamiltonian that acts linearly on that state. Linear quantum evolution does not mean that particles cannot interact. Calling particles patterns therefore does not decide whether they pass through, scatter or form a bound state.

Identical-particle statistics matter as well. Exchanged alternatives can interfere with the sign required for fermions or bosons. The effect is

included in the amplitude; it is not a separate decision taken by particles when they recognise one another.

The most useful questions are specific. What was prepared? Which final states can the detector distinguish? What interaction connects the states? Which approximation controls the calculation? Those questions turn a verbal image of two things meeting into a prediction we can compare with counts.

Try 7. Why does passing an energy-threshold check not by itself predict how often a reaction occurs?

A balanced charge and energy account is necessary, but it does not supply a rule for how often a reaction happens. A physical theory must also specify the interactions. In field theory these rules are collected in a quantity called the action, from which the evolution equations are obtained.

8 Which interactions are allowed?

Before reading on. If a proposed interaction conserves electric charge, does that make it a complete acceptable term in a physical theory?

Charge conservation is necessary for many familiar processes, but it is far from the only condition. We must identify the fields, their transformation properties, the spacetime symmetries and the other charges. An interaction must fit the whole specified theory.

In a field action, a vertex means an interaction term. In a Feynman diagram, it is where the corresponding lines meet. It is not a geometrical corner of a spatial cell. Confusing those meanings can make a drawing look more explanatory than it is.

Why do the conservation conditions hold? In the usual flat-spacetime field theory, an action unchanged by shifts of time and position has corresponding conserved total energy and momentum. Internal symmetries constrain the charge balance of its interaction terms. The diagram inherits these properties from that action. The dot where lines meet does not enforce them independently. A general curved, evolving spacetime

requires more care with global energy; the local gravitational account is discussed in the companion books.

Suppose a reaction has total charge $+1$ before the interaction. Final charges $+1$, $+1$ and -1 also total $+1$. The charge ledger balances. But the reaction may still be forbidden by energy, angular momentum or another conserved quantity. Even an allowed reaction needs a nonzero coupling and an amplitude.

Gauge symmetry concerns changes of the internal mathematical description that leave physical predictions unchanged. Requiring consistency under these changes strongly constrains how charged fields interact with gauge fields. It does not fix every coupling, mass or mixing parameter. Nor does it generally leave only one possible action. Effective field theories include additional symmetry-allowed terms, organised by the energy scale at which their effects matter.

A restriction to a particular short list of terms can be physically motivated, but that restriction must be declared. Locality, unitarity and consistency under changes of description impose further demands. None is supplied just by giving an operator a plausible name.

This offers a fair test for an algebraic proposal. An algebra may provide a precise space of possible operations. A physical principle would then need to select which operations occur and constrain their coefficients. If those choices are inserted to match the known theory, the construction is a representation of that theory. It may still be valuable, but its explanatory claim must match what was actually selected.

Try 8. The charges $+1$, $+1$ and -1 sum to what? Name one further check needed before declaring a reaction possible.

An interaction can remove particles of one kind and produce others while preserving the total conserved quantities. Electron-positron annihilation is a familiar example. It also raises a larger question: why did cosmic history leave so much more ordinary matter than antimatter?

9 What is antimatter, and why did matter survive?

Before reading on. If an electron and positron annihilate, does their rest energy disappear?

The positron is the electron's antiparticle. It has the same mass and opposite electric charge. More generally, antiparticles have corresponding opposite additive charges; some neutral particles can be their own antiparticles. Antimatter is not matter with negative mass or negative energy.

An electron-positron pair can annihilate into radiation, subject to conservation laws. In their centre-of-momentum frame, annihilation into two photons sends the photons in opposite directions with equal energies. The initial rest and kinetic energy becomes energy of the outgoing state. Nothing has been deleted from the energy account.

For a simple example, take 2 units of total incoming energy and zero total momentum. In the two-photon final state, each photon carries 1 unit. Their opposite momenta sum to zero. Trying to put all that energy into one photon in empty space would fail the momentum condition.

Why the observable universe contains a net excess of ordinary matter over antimatter is a separate question. It concerns cosmic history, particle interactions and the generation or survival of an asymmetry. The existence of antiparticles does not require the universe to contain equal present amounts of each.

Baryon number counts baryons, such as protons and neutrons, minus antibaryons; a quark carries one third of a baryon's count. A baryon asymmetry is a net excess of one over the other. Under the usual assumptions, generating an initially absent baryon asymmetry requires appropriate baryon-number violation, asymmetry under charge conju-

gation and under CP, and departure from thermal equilibrium. Charge conjugation exchanges particles with their antiparticles. P denotes a mirror reversal; CP combines the two transformations. Equilibrium means a state with steady thermal populations rather than a continually maintained imbalance. These are the familiar Sakharov conditions; they identify necessary ingredients within that setting, not a unique mechanism or an abundance calculation.

Several proposed histories attempt to supply the ingredients. Their rates, initial conditions and later evolution must reproduce the observed asymmetry and survive other tests. The residual excess was small compared with the early thermal populations, yet it supplied the ordinary matter from which stars and planets later formed.

Try 9. Why can a pair initially at rest as a whole annihilate into two opposite photons, while a single photon would fail the same momentum account?

Part two

What the weak interaction changes

A free neutron can turn into a proton while producing an electron and an electron antineutrino. This is beta decay, one of the processes described by the weak interaction. The particles after the event differ from those before it, yet total electric charge, energy and momentum must still balance.

The weak interaction also distinguishes a system from its mirror image. This makes the word handedness relevant, although a spinning corkscrew will turn out to be an unreliable picture. We need to distinguish the direction of spin relative to motion from a different property of the field description.

The W and Z fields enter the theory of these processes. Their particle excitations are heavy and short-lived; ordinary neutron decay has far too little energy to produce a real W. That is not a contradiction. The following questions connect the symmetry of the interaction, the charge account, and the difference between a field's internal contribution and a particle seen in a detector.

10 Why does the weak interaction prefer one handedness?

Before reading on. If an observer overtakes a massive particle, can the particle's direction of motion reverse in that observer's view without its spin reversing?

Yes. This is one reason that handedness needs a careful definition. Helicity compares spin with the direction of motion. A sufficiently fast observer can reverse the momentum direction of a massive particle by overtaking it, changing that comparison.

Chirality refers to a different split of the field into two kinds of components, called left and right. They transform in different ways under relativistic changes of frame. Defining that split needs the spinor mathematics; it is not obtained by inspecting a spiral or by changing the observer's speed. In the Standard Model, a charged weak current is the field combination through which a W changes one fermion type into another. These currents couple to left-chiral fermion fields and their corresponding antiparticle components. The left and right field components are not treated symmetrically. This is the precise content behind the phrase left-handed weak interaction.

For very relativistic particles, helicity is often a useful way to discuss the observed effect, because the relation to chirality becomes simple in the negligible-mass limit. For massive particles, however, the two words cannot be exchanged in every sentence. A mass term links the chiral components.

The asymmetry is experimentally important. A mirror reversal changes spatial handedness, and weak processes need not have the same probabilities as their mirror-reflected versions. Parity violation was a discovery about physical laws, not an artefact of choosing a left-handed coordinate system.

Why this particular representation pattern occurs remains a deeper question. The Standard Model incorporates it and calculates its consequences. A microscopic explanation would have to recover both the field content and the interactions, including the consistency conditions of the gauge theory.

A spiral drawing by itself does not do that. A clockwise-looking curve can change appearance when viewed from the other end, whereas the physical coupling must be defined independently of our sketch. The useful test is which prepared states interact and with what amplitudes.

Try 10. Which word compares spin with momentum: helicity or chirality? Why is overtaking relevant to it?

The handedness of the weak interaction concerns the states it couples to. Electric charge poses a separate bookkeeping requirement. In beta decay the quark changes type, but the particles produced elsewhere in the same event must keep the total charge unchanged.

11 Why does a charged weak process change charge by one unit?

Before reading on. A down quark becomes an up quark, changing from charge $-1/3$ to $+2/3$. Has electric charge been violated?

The quark's charge has increased by one unit. The complete process must include another part carrying the compensating charge. In a charged-current weak description, a W field connects the appropriate members of a weak multiplet: a set of field components grouped together by the weak symmetry, such as up-type and down-type quarks. The charged W bosons carry one positive or negative unit.

For beta decay at the quark level, the useful charge account is a down quark becoming an up quark while a negative unit is passed to the lepton part of the process. The electron carries that negative charge; the electron antineutrino is electrically neutral.




Figure 3: The charge account for neutron beta decay. The full initial and final systems both have zero electric charge. The diagram records the reaction; it does not put a real, freely propagating W boson inside the neutron.

At the hadron level, a neutron of charge zero becomes a proton of charge $+1$, an electron of charge -1 and an antineutrino of charge zero. The total is $1 - 1 + 0 = 0$. The weak interaction changes particle identities while preserving total electric charge.

The one-unit difference follows from the electroweak representation structure and the charged currents in the specified theory. It is not a statement that every weak event changes a particle's electric charge. Neutral-current weak processes involve the Z field and can leave the external particle's charge unchanged.

Nor does balancing electric charge alone determine the event rate. The weak coupling, mixing, available energy and the initial and final states all enter. In nuclear beta decay, the relevant nuclear transition can strongly affect the probability.

This is a good example of the difference between a ledger and a dynamics. The ledger is essential: an unbalanced charge account exposes an error immediately. Many processes can share a balanced ledger, however, while having very different amplitudes.

Try 11. Write the total final electric charge in neutron beta decay using the three numbers in the text.

The charge account for beta decay mentions the W field. A real W particle would require far more energy than a neutron decay releases. To understand the calculation we must distinguish that particle production threshold from the field's contribution to a low-energy interaction.

12 Does beta decay briefly make a lighter W ?

Before reading on. A decay has far too little available energy to make a real W boson. Can the W field still contribute to the interaction?

Yes. A field's contribution to an amplitude is not restricted to the production of a freely propagating, detectable particle of that field. At energies well below the W mass, charged weak processes can be described by an effective interaction among the lighter fields.

One way to obtain that description is to start with the W propagator, the mathematical contribution connecting the two interactions in the full theory, and expand it for small momentum transfer relative to the W mass scale. Its leading effect becomes a local interaction in the low-energy approximation. The heavy field has left a calculable effect even though a real W cannot be produced.

The internal W line is often called virtual. It does not describe an on-shell W whose rest mass temporarily shrinks to fit the energy budget. On shell means satisfying the physical free-particle energy–momentum relation. The internal contribution is part of the amplitude. Its transferred momentum can be fixed by the external particles; other internal variables are integrated over when the diagram requires that operation.

Nor is there a loan of the missing rest energy that uncertainty allows for a short interval. The decay obeys the conservation laws of the full process. The final proton, electron and antineutrino share the energy actually available. No detector has to find a real heavy W in between.

The approximation has a domain of validity. As the energy approaches the electroweak scale, the simplified low-energy description must be replaced by the full propagator and the relevant theory. A successful approximation at low energy is not a claim about all collisions.

This distinction is familiar in other areas of physics: a degree of freedom can affect a low-energy response without being excited as a real outgoing object. Quantum field theory gives a systematic way to calculate that effect and its corrections.

Try 12. Which description applies to the W in ordinary beta decay: an observed lighter W , or an internal field contribution whose low-energy effect can be approximated?

At a high-energy collider, a heavy unstable excitation can be produced and inferred from its decay products. A concentration of events around a characteristic reconstructed energy is called a resonance. Its position and width carry information, but neither identifies all the particle's properties by itself.

13 Is a short-lived particle just energy at a frequency?

Before reading on. If two excitations have the same energy, must they be the same kind of particle?

Energy is one property of a state. It does not specify electric charge, spin, colour or the interactions that prepare and detect the state. A photon and a massive particle can have the same total energy while behaving very differently.

The relation between energy and phase frequency is fundamental in quantum evolution. It says how the phase of an energy component changes with time. It does not identify a particle species from a frequency alone. A physical identification requires the other quantum numbers and the response laws.

W and Z bosons are short-lived massive excitations with specific electroweak properties. Their production and decay show up as resonant structures in suitable collision data. A resonance is a concentration of events associated with a characteristic energy and a finite width.




Figure 4: A schematic resonance curve. The horizontal axis is reconstructed energy; the vertical axis represents event rate. The shape illustrates a width, not a fit to W or Z data.

Under the standard resonance description, a shorter lifetime corresponds to a broader natural energy width. Doubling the decay rate halves the lifetime and doubles that width. Detector resolution and the production process can also broaden an observed distribution, so the measured shape needs analysis.

This does not mean a short-lived particle has an arbitrary mass chosen anew in each event. An unstable state is characterised through its resonance and couplings; the precise mass and width conventions must be stated in precision work.

The useful experimental question is whether one consistent description accounts for the resonance, spin, charges, production rates and decay patterns. A wavy drawing with the right frequency answers only a small part of that question.

Try 13. In the ideal lifetime–width relation, what happens to the natural width when the lifetime is halved?

The low-energy effect of a heavy mediator is suppressed, while production near a resonance can be much more prominent. The word weak therefore needs a specified process and energy range. We can test that point with a simple comparison that holds the other inputs fixed.

14 Is the weak interaction always weak?

Before reading on. If a process is strongly suppressed at low energy, must it remain equally suppressed when the collision energy approaches the mediator’s mass scale?

Weak names a particular interaction. It is not a ranking that applies unchanged to every process at every energy. At energies far below the W and Z masses, many weak amplitudes are suppressed by the heavy mediators. That helps explain the long times and small interaction probabilities familiar from low-energy weak processes.

The low-energy effective interaction contains an inverse heavy-mass scale. Schematically, keeping the underlying coupling fixed, doubling the mediator mass reduces the leading coefficient by a factor of four. A rate proportional to the square of that coefficient would then be reduced

by a factor of sixteen if the other factors describing the available energies, momenta and final states were unchanged. This is a comparison within a stated approximation, not a formula for every weak rate.

As the collision energy rises, those kinematic factors change and the low-energy approximation eventually fails. Near a resonance, real W or Z production is possible in appropriate channels. The full electroweak theory describes the transition.

Comparing an electromagnetic rate with a weak rate therefore requires matched physical questions. The incoming states, energy, accessible final states and couplings all matter. A charged particle may have a large electromagnetic response in one experiment while a neutral particle has no corresponding channel at all.

Symmetries and conservation rules can also suppress a process independently of the nominal strength of its interaction. A forbidden strong decay has zero rate even though the strong interaction is not generally small. An allowed weak decay may then dominate.

The label is useful once we know what it refers to. For predictions, replace this force is weaker with an actual process and the energy range in which the comparison is being made.

Try 14. In the stated low-energy scaling example, doubling the mediator mass changes the coefficient and its squared contribution by what factors?

Part three

Following neutrinos

Neutrinos are produced in beta processes, in the Sun and in many other reactions. A detector usually learns about one only when a rare interaction produces something easier to record, such as a charged particle or a changed nucleus. The absence of a click normally means that the neutrino passed through without a detectable interaction.

The word flavour names the electron, muon and tau neutrino alternatives selected by weak production and detection. It does not describe a shape or a substance carried inside the particle. The combinations selected by these weak processes differ from the combinations that propagate with definite masses. This mismatch makes interference during flight observable.

We start with detection, then follow the phase changes that produce oscillations. We distinguish what oscillation measurements tell us about masses from what they leave unknown. The final example returns to a concrete detector: a neutrino capture that changes a chlorine nucleus into argon. Energy and charge accounting will matter as much as the names of the particles.

15 Why are neutrinos hard to detect?

Before reading on. Does a particle's small mass, by itself, imply a small probability of interacting in a detector?

The neutrinos of the Standard Model carry neither electric charge nor QCD colour. They therefore lack the ordinary electromagnetic and strong interactions that make many other particles easy to stop. Their weak interactions have small probabilities at familiar energies. Small mass is a separate fact.

A detector compensates by offering many targets, collecting many incoming neutrinos and waiting. Most neutrinos leave no record. A rare interaction can create a charged particle or excite a nucleus, producing a signal that the apparatus can amplify and identify.

Here is a counting example. Suppose the chance of a detectable interaction per crossing neutrino is one in a million, under specified conditions. Ten million crossings then give an expected ten events. Expected does not mean exactly ten in every run. Backgrounds, efficiency and statistical fluctuations must be included before the count can be interpreted.

Real probabilities depend on neutrino energy, target material, interaction channel and detector design. There is no universal one-in-a-million number for neutrinos. The example only shows how a small single-particle chance can become measurable with a large flux and target.

Different detector types answer different questions. Some measure light produced by charged particles. Others collect ionisation or count a changed nuclear species. An interaction can be physically possible yet fall below a particular instrument's threshold.

This is why a detector record cannot be equated with the entire neutrino state. It contains the information that the chosen interaction and apparatus made available. Reconstructing energy, direction or flavour requires a model of that measurement process and its limitations.

Try 15. In the illustrative example, what is the expected event count for twenty million crossings at the same interaction probability?

A neutrino detector can sometimes identify which charged lepton was produced, and hence which flavour was tested. Experiments compare that result with the flavour selected at production. Their difference suggests that propagation involves more than a fixed catalogue label.

16 Why do neutrinos change flavour in flight?

Before reading on. Could flavour probabilities change even if the masses of the propagating components remain fixed?

The weak interaction defines neutrino flavours by the charged leptons involved in production or detection. The states that propagate with definite masses are different combinations. A prepared flavour can therefore contain several mass components.

Those components acquire phase at different rates. When a later weak measurement tests flavour, the components combine again. Their changed relative phases alter the flavour probabilities. The masses have not changed during the journey. This is interference in a different physical setting.




Figure 5: An idealised two-component account. The branches label propagation components, not visibly separated paths. Production and detection test flavour combinations.

Use the balanced interferometer arithmetic from the first volume as a teaching model. Prepare an equal combination of two propagation states. At zero relative phase, the original flavour has probability 1. At a quarter-turn, the two flavour probabilities are each one half. At a half-turn, the other flavour has probability 1. A full turn restores the starting prediction.

This equal-mixture two-flavour example is deliberately simple. Actual neutrino oscillations require three-flavour mixing parameters, and propagation through matter can change the effective evolution. Wave packets must remain sufficiently coherent for the interference description to apply.

Why not electrons turning into muons along an ordinary wire? The charged leptons ordinarily prepared and detected are mass eigenstates. Neutrino production selects coherent flavour combinations of distinct mass eigenstates. The preparation, phase evolution and detection basis together produce the oscillation pattern; changing a catalogue label does not.

For highly relativistic neutrinos in vacuum, the leading oscillation phase depends on mass-squared differences and the ratio of travel distance to energy. This is why oscillations give information about differences without directly fixing the absolute mass scale. The optional pages show that arithmetic separately.

Try 16. In the equal-mixture example, what happens to flavour probabilities after a half-turn of relative phase? Have either of the masses changed?

Oscillations reveal a phase difference accumulated during flight. That does not immediately give each neutrino mass separately. We must distinguish the information in differences from an absolute mass scale, and separate both from the question of particle-antiparticle identity.

17 Why are neutrinos so light, and are they their own antiparticles?

Before reading on. If oscillations reveal differences between squared masses, do they determine every mass separately?

Differences do not fix an overall level. Three squared masses could, as a purely arithmetic example, be 1, 4 and 10 in chosen units. Adding 5 to all three gives 6, 9 and 15. The pairwise differences are unchanged. The example is not a set of measured neutrino masses.

Oscillation experiments constrain the mass-squared differences and mixing. Nonzero independent differences require at least two massive states. They do not, by themselves, require the lightest mass to be nonzero. Other measurements and cosmological inference probe the absolute scale with their own assumptions.

The original minimal Standard Model has massless neutrinos. Describing their masses requires additional structure. Possible mechanisms include very small couplings or effects associated with heavy new degrees of freedom. A seesaw mechanism introduces mixing between light and very heavy neutral fields. In a suitable regime, increasing the heavy mass at fixed mixing terms can make the light mass smaller, but that explanation is conditional on the extra fields and couplings it assumes.

There is a separate distinction between Dirac and Majorana mass. A Dirac neutrino has a distinct antiparticle. A Majorana neutrino is its own antiparticle in the field-theoretic sense. Zero electric charge removes one obstacle to the second option. Other exact conserved charges would matter too: a Majorana mass cannot preserve a lepton-number distinction between that particle and its antiparticle. Lepton number is the proposed count that assigns opposite signs to leptons and antileptons. Electrical neutrality alone does not prove that nature has chosen a Majorana mass.

Weak measurements can distinguish helicity-dependent behaviour even in a Majorana description. Its own antiparticle must not be translated into every neutrino measurement gives the same result as every antineutrino measurement.

In ordinary double-beta decay, a nucleus changes two neutrons into protons while emitting two electrons and two antineutrinos. Neutrinoless double-beta decay would emit the two electrons without those antineutrinos. Searches for it probe lepton-number violation and possible Majorana physics. Interpreting a signal or a null result requires the assumed mechanism and nuclear response. Oscillations alone do not answer this question, and an explanation of small mass must be assessed alongside those additional tests.

Try 17. What are the successive differences in the two illustrative lists of squared masses? Why can neither list be selected from those differences alone?

One pioneering solar-neutrino method counted argon atoms produced in a chlorine target. Here the detector record is a changed nuclear species. Understanding the record requires an ordinary reaction account as well as the neutrino's quantum description.

18 Where does the energy come from when a neutrino changes chlorine into argon?

Before reading on. If a neutrino induces a nuclear reaction, may we leave its incoming energy out of the balance?

A solar electron neutrino can be captured in a charged-current reaction on chlorine-37, producing argon-37 and an electron. Within the nucleus, the weak process changes a neutron into a proton. The proton number increases by one, so the element changes.

The incoming neutrino supplies energy and momentum. Whether the reaction can occur depends on the whole initial and final states: the participating nuclei, the emitted electron, recoil and any nuclear excitation. There is a threshold, not a permission to borrow energy from the vacuum.

The charge account is easy. A chlorine nucleus has 17 positive units. An argon nucleus has 18, and the outgoing electron carries -1. Thus $18 - 1 = 17$, matching the initial nuclear charge. Atomic electrons and their binding must also be handled consistently when using atomic masses for a precise threshold.

Imagine, as a deliberately invented threshold example, that a reaction requires 2 units of incident energy. A neutrino bringing 3 units can leave 1 unit to be distributed among the permitted final motion and excitation, once the threshold cost is paid. A neutrino bringing only 1 unit cannot trigger that channel. These numbers are not the chlorine threshold.

A radiochemical detector can count the newly produced argon atoms. Its result is an accumulated record of captures, rather than a photograph of each incoming neutrino's trajectory. Extracting a solar-neutrino flux from the count requires the reaction response, extraction efficiency, exposure and backgrounds.

The example connects three layers of description. The weak interaction permits the identity change. Nuclear energies set which incoming neutrinos can cause it. The apparatus turns some reactions into countable records. All three are needed to interpret the experiment.

Try 18. Why is the final nuclear charge 18 compatible with an initial nuclear charge 17?

Part four

What makes something stable?

A particle or nucleus is called stable against a particular decay if that decay cannot occur under the stated conditions. A long lifetime can have several explanations. There may be no accessible final state, a conservation law may prohibit the change, or the interaction may allow it only very slowly. A barrier is one possibility, not a universal explanation.

For a nucleus the relevant object is the entire bound system. Protons repel electrically, while the short-range nuclear interaction binds protons and neutrons. Changing one neutron into a proton changes both the particle count and the binding energy. We must compare complete initial and final systems before deciding whether a decay is possible.

In this part, an energy diagram is an account of those whole-system energies. It is not a picture of a particle climbing a literal hill. The questions move from simple restrictions to nuclear examples, asking first whether a channel is open and then what could determine its rate.

19 Does stability require a high barrier?

Before reading on. A particle has no allowed lower-energy final state with the required conserved charges. Does it also need a high barrier to prevent decay?

A barrier is one way to make a state long-lived, but it is not the only way. An isolated state cannot decay if no accessible final state satisfies the theory's conservation rules. There may be no permitted destination at all.

The electron is the lightest known particle with its electric charge. Within the usual theory, electric charge conservation and the available spectrum prevent its decay into lighter states. This argument is different from imagining an electron trapped behind a large energy hill.

For a metastable state, by contrast, a lower-energy configuration exists but the transition is slow. Tunnelling through a barrier or a small interaction amplitude can produce a long lifetime. Radioactive nuclei provide familiar examples of long-lived states that are not absolutely stable.

Stable also needs an environment. A state stable in isolation may react when struck by another particle. A bound atom can be ionised by sufficient radiation; a proton can participate in nuclear reactions without undergoing spontaneous proton decay. Stability does not mean unresponsiveness.

Consider an invented energy diagram. An initial state has 10 units. One charge-conserving final configuration needs 12 and is inaccessible in isolation. Another would need only 8 but has the wrong total charge. Neither is an allowed decay destination. The low energy of the second does not rescue its forbidden charge change.

A lifetime claim should therefore identify the theory, the available channels and the physical conditions. An experimental lower bound means no decay has been established at that sensitivity. It is not automatically a proof of eternal survival in every possible extension of the theory.

Try 19. Why are both invented final configurations ruled out, for different reasons?

The electron example showed how a missing decay destination can protect a particle. The opposite case is a state with a readily accessible channel. Comparing the top quark with the proton makes clear why mass alone cannot rank lifetimes.

20 Why does the top quark decay rapidly while the proton lasts?

Before reading on. Does a larger mass always make a particle more stable?

Mass affects the energy available to a decay, but it does not determine stability by itself. A heavy particle can have many accessible lighter final states. A light one may have none consistent with its conserved charges.

The top quark has a rapid allowed weak decay, predominantly to a bottom quark and a W boson. Its large mass opens this channel. The relevant coupling and final-state phase space make the decay fast. Phase space means the set of momenta and states allowed by the conservation conditions.

The proton has a very different position in the spectrum and charge structure. No proton decay has been established. Common low-energy Standard Model processes preserve baryon number, while many extensions can permit proton decay with strongly suppressed rates. Experimental limits test particular decay channels. Calling the proton immortal goes beyond those observations.

A simple comparison shows why mass alone is insufficient. Imagine two states, each with 10 units of rest energy. For one, an allowed final state costs 6 units and couples strongly. For the other, every lighter state has an incompatible conserved charge. Equal initial energy does not imply equal lifetime.

When several decay channels are available, their partial rates add to give the total decay rate. The fraction entering a particular channel is its branching fraction. Measuring those fractions tests the interaction more sharply than a single statement that the particle is short-lived.

For any proposed model, the useful calculation is a decay rate or a selection rule with its assumptions stated. That gives experiments something to test. The adjective robust does not supply the rate.

Try 20. Name the three ingredients, beyond merely stating a mass, that the text uses to explain a decay rate.

A free neutron decays, yet many nuclei containing neutrons persist. Binding has changed the system whose energy must be compared. We must ask about the parent nucleus and the complete daughter state, not give each neutron the decay opportunities of an isolated one.

21 Why can a bound neutron outlive a free neutron?

Before reading on. If an isolated neutron can beta-decay, must every neutron inside a nucleus have the same decay opportunity?

A free neutron can decay into a proton, electron and electron antineutrino. The available final state has lower rest energy, and the weak interaction permits the transition. The released energy is shared among the final particles' motion.

The mean lifetime of a free neutron is about fifteen minutes, not several hours; a mean lifetime is not a deadline for every individual neutron. A sample of free neutrons gains protons through this decay. That does not imply that every neutron in the universe steadily disappears: binding and environmental reactions change the available processes.

Inside a nucleus, we must compare whole systems. Turning one neutron into a proton changes the nuclear binding, electric repulsion and allowed quantum states. The daughter nucleus may be heavier than the parent once the complete reaction is accounted for. Then that decay is energetically closed.




Figure 6: An illustrative whole-system comparison. Only a lower accessible final energy can receive the released energy in an isolated decay. These levels are teaching values, not measured nuclear masses.

For example, assign the parent 100 units of total energy. If the complete final system requires 101 units, the decay cannot proceed in isolation. If another parent also has 100 units but its final system requires 98, there are 2 units available for final motion, subject to the other rules. The word neutron alone does not decide between these cases.

This is why some neutron-rich nuclei beta-decay and others are stable against that channel. It is not that the bound neutron has acquired a universal exemption from the weak interaction. Its many-body environment changes the accessible states.

Using tabulated masses requires care about whether they refer to nuclei or neutral atoms, and about the final electron and any excited daughter state. An inconsistent choice can manufacture or remove a threshold on paper.

Even when beta decay is energetically permitted, the nuclear transition matrix element and selection rules affect its rate. The matrix element is the amplitude associated with the interaction acting between the specified initial and final nuclear states; it measures more than whether their energies permit a change. An energy balance opens the door; it does not tell us how quickly the system passes through it.

Try 21. In the 100-to-98 example, how much energy remains for allowed final motion? What still has to be checked before assigning a rate?

Changing the neutron-to-proton balance changes nuclear binding and electric repulsion together. Isotopes are nuclei with the same number of protons but different numbers of neutrons. Their different stabilities let us test how those contributions compete.

22 Why do heavy nuclei need extra neutrons?

Before reading on. If neutrons add nuclear attraction without adding electric charge, should adding more and more neutrons always make a nucleus safer?

Protons repel each other electrically. The residual strong nuclear interaction binds protons and neutrons over short ranges. As proton number grows, electric repulsion becomes a larger part of the energy balance. Additional neutrons can help binding without adding electric repulsion.

That is only one side of the account. Neutrons are fermions, as are protons. Filling more neutron states costs energy, and the nuclear interaction also penalises large neutron–proton imbalance. Too many neutrons can favour beta decay or become unbound.

The stable region is therefore a compromise among several contributions to the energy of the entire nucleus. For lighter nuclei, similar proton and neutron numbers are often favoured. For heavier ones, the stable region shifts towards more neutrons. It is a trend with structure, not a universal simple ratio.

Pairing and shell effects create additional structure. Some particle numbers complete particularly favourable arrangements of nuclear states. A smooth liquid-drop account captures broad trends but needs these quantum corrections to describe individual nuclei accurately.

Very heavy systems also have competing routes to lower energy, including alpha decay, which emits a helium nucleus, and fission, which splits the nucleus into substantial fragments. A fission barrier can make an otherwise energetically favourable split slow. This is where barriers and tunnelling return to the stability discussion.

An island of stability refers to predicted regions where shell effects may make some superheavy nuclei longer lived than their neighbours.

It does not promise absolute stability or an unlimited sequence of stable elements. Different decay channels must still be assessed; lead and uranium do not define one universal boundary that settles every possible isotope.

As a counting illustration, a nucleus with 40 protons and 50 neutrons has 90 nucleons and 40 positive charge units. Adding two neutrons changes the nucleon count to 92 while leaving electric charge unchanged. That count alone cannot tell us whether the new isotope is more stable: its total energy and decay channels must be calculated.

This is the common lesson of the volume. The catalogue and the ledger constrain the question. Dynamics and the available states determine the answer.

Try 22. In the illustrative nucleus, what changes when two neutrons are added, and what cannot be inferred from that count alone?

Part five

Questions behind the familiar answers

We have described particles through their charges, masses and allowed interactions. The Standard Model brings those descriptions into one quantum field theory. It gives many accurate predictions, but it does not make every ingredient inevitable.

The following questions distinguish three sorts of answer. Some apparent puzzles come from a misleading picture, such as a positron literally visiting yesterday. Some concern mechanisms that the theory describes, such as confinement and the Higgs mechanism. Others ask why a permitted parameter has its observed value. A theory can succeed at the first two tasks while leaving the third open.

No new algebra is needed here. When a symmetry or a field appears, we ask what it changes in a possible experiment and what it leaves for separate measurement.

23 Why can we not pull one quark out of a proton?

Before reading on. If increasing the separation of two constituents stores more energy in their field, must supplying enough energy eventually release the same two constituents as isolated particles?

That conclusion would be reasonable for some familiar bound systems. An electron can be removed from an atom by supplying its ionisation energy. Quarks in ordinary hadrons behave differently.

In quantum chromodynamics, the strong field interacts with itself. For separated colour sources, the field does not simply spread and weaken like an electrostatic field. In an appropriate separation range, calculations describe energy concentrated in a connecting region, often called a flux tube. Pulling the sources farther apart then costs increasing energy.

When the quantum theory allows light quark–antiquark pairs to be created, a different outcome can become possible: the field energy can produce a quark–antiquark pair. The original and newly produced constituents reorganise into colour-singlet hadrons. Continuing to supply energy can give more hadrons instead of an isolated original quark. The phrase string breaking describes this energy redistribution; it does not mean a material string has been found inside the proton.

For a deliberately simple accounting example, suppose a configuration contains five extra energy units and creating an allowed pair costs three. Two units remain for motion and other excitations. This is only a threshold illustration. Actual production requires the interaction amplitude, momentum balance and the full quantum state.

High-energy collisions do resolve quark and gluon behaviour. Their outgoing energy appears in jets of hadrons, which carry information about the short-distance process. A jet is not a freely propagating coloured quark arriving intact at a distant detector.

Confinement is supported by the observed behaviour of hadrons and by calculations such as lattice QCD. These calculations can treat strong interactions without assuming that a few terms of a weak-coupling expansion suffice; this is what nonperturbative means here. A picture of three phase arrows cancelling does not establish that dynamics. Nor is every state of strongly interacting matter a collection of separate protons and neutrons: hot, dense QCD matter has a different description.

It is also too restrictive to say that only ordinary three-quark and quark–antiquark hadrons can occur. More complicated colour-singlet combinations are allowed. The elementary neutrality rule and the dynamical question of which combinations bind are different questions.

Try 23. In the five-unit accounting example, how much energy remains after a three-unit pair-production cost? Why does that subtraction not calculate a production rate?

Quark pair production required new particles in the final state. Antiparticles also appear as lines with reversed arrows in diagrams. That drawing convention raises a different puzzle about time.

24 Can a positron really travel backwards in time?

Before reading on. If a collision produces a positron that is detected later, does a backward arrow in a calculation imply a detector could have recorded it before the collision?

The observed positron carries positive energy and travels from its production event towards a later detection or annihilation event. It has the electron's mass and opposite electric charge. Its track does not require it to emerge in the laboratory's past.

The backward-time description belongs to the organisation of a relativistic quantum calculation. A fermion propagator joins two interaction points and incorporates both particle and antiparticle contributions. One part can be represented as an electron contribution with the time ordering reversed, or as the corresponding positron contribution with the physical time ordering used for antiparticles. These are related ways of arranging an amplitude.

An arrow on a fermion line therefore keeps track of charge flow and the algebraic convention. It is not automatically the direction in which a clock attached to an observed particle runs. The time-ordered propagator itself is also not a direct photograph of a particle's path between measurements.

For a concrete energy account, let an electron and a positron be created with opposite momenta. The incoming state must supply their combined rest energy and any kinetic energy. The final charges add to zero. Describing the positron line differently cannot remove those requirements or make the creation event unnecessary.

Relativistic field theory is built to preserve causal restrictions on observable signalling. Some intermediate mathematical functions have properties that would be misleading if read as ordinary trajectories;

observable probabilities require the complete calculation. A diagram by itself does not authorise sending a message into the past.

Feynman's description is valuable because it exposes a connection between relativistic amplitudes and antiparticles. It becomes misleading when every line in the drawing is treated as a miniature film of a real object. The useful question is what the representation predicts for prepared states and detected outcomes.

Try 24. Which must remain unchanged when a positron contribution is redrawn: the measured charge and energy balance, or the direction of every arrow on the page?

Antiparticle arrows are one example of language that compresses a calculation. Another is the claim that a massive force carrier has eaten part of the Higgs field. Here the missing explanation is about the independent ways a field can be excited.

25 What does it mean for the W and Z to eat a field?

Before reading on. If a reformulation combines degrees of freedom into new particles, have those degrees of freedom been destroyed?

A degree of freedom is an independent part of a physical description that can vary. For a freely travelling photon there are two physical polarisation states. A massive spin-one particle has three spin states. The additional one is often called longitudinal: its polarisation is associated with the direction of motion.

The Higgs field can be described by four ordinary real-valued components, rather than four separate kinds of Higgs particle. They are scalar components in the electroweak theory. Scalar here refers to their behaviour under spacetime rotations, not to an absence of internal electroweak labels. Around the chosen vacuum, the field and gauge excitations have to be described together.

In the usual perturbative account, which studies small field excitations around that background, three scalar directions supply the additional

longitudinal degrees of freedom of the two charged W states and the neutral Z. One independent scalar excitation remains as the Higgs boson. The photon keeps its two transverse polarisations, whose electric-field oscillations are perpendicular to its direction of travel. The physical degrees of freedom have been reorganised.




Figure 7: A count of physical perturbative degrees of freedom in the electroweak Higgs mechanism. It is not a sequence of particles swallowing other particles, and not a model of the thermal history of the early universe.

Counting alone does not produce this result. The coupled field equations, the Higgs potential and the chosen vacuum determine the mass terms and interactions. Gauge freedom lets the same physical content be represented with different field variables. In one convenient representation, the three scalar directions no longer appear as separate physical particles. That is the origin of the eating metaphor.

It is not two ordinary waves cancelling out of phase. Such cancellation changes an amplitude at an output; it does not explain a new longitudinal polarisation or the characteristic couplings of a massive gauge field. Nor does the mechanism predict the electron’s or muon’s mass without their separate coupling parameters.

The Standard Model’s account is tested through the W, Z and Higgs properties and their interactions. A different microscopic proposal would have to recover that behaviour, not merely count four labels and discard three.[6]

Try 25. Count the degrees of freedom in the diagram: eight gauge polarisations plus four scalar components before reorganisation; how many appear in the final description?

The Higgs mechanism used a particular electroweak gauge theory. Colour used another symmetry structure. We can now ask why those specific structures occur in the Standard Model.

26 Why these three gauge symmetries?

Before reading on. If an encoding has three distinguishable internal labels, does it already follow that its interactions obey colour $SU(3)$?

Three labels specify a set of possibilities. A symmetry specifies allowed transformations of states, and a gauge theory specifies how a local change of mathematical description leaves observable predictions unchanged. These are additional structures.

In the familiar notation, the Standard Model uses colour $SU(3)$, weak $SU(2)$ and hypercharge $U(1)$. At this level the notation names three families of internal transformations. The quark and lepton fields occupy particular representations, meaning specified ways of changing under those transformations.

The usual local description has eight colour gauge fields, three weak ones and one hypercharge field. Eight plus three plus one makes twelve. Following electroweak symmetry breaking, the weak and hypercharge fields are reorganised into the two charged W states, the Z and the photon. The count is not a prediction of the numbers of matter generations.

These choices are constrained by experiment and consistency. Quantum anomalies, for example, can spoil a gauge symmetry that appears valid in the classical equations. The Standard Model's matter assignments cancel the relevant gauge anomalies. That is a serious restriction, but it does not by itself prove that no other consistent theory can exist.

Gauge symmetry also leaves coefficients to be supplied. Different couplings can respect the same symmetry. To predict their values, a proposed foundation needs an additional physical principle or dynamics that selects them, followed by a comparison with measurements.

These three group names specify the transformations needed for the local discussion here. A full gauge theory also has to say how those descriptions fit together over larger regions and which charge patterns

are allowed there. We do not need that further mathematical machinery to see why the group names alone cannot fix interaction strengths.

The same caution applies to a qubit register. A useful matching of its labels with particle labels can organise a model. To derive the interaction theory, one must also obtain the action, representations, locality, constraints and coefficients. Counting the labels alone cannot supply those missing choices.

Try 26. The local gauge-field count is eight plus three plus one. Does reproducing that total determine their interaction strengths?

A specified symmetry can allow a term whose coefficient is then surprisingly small. The strong interaction supplies an important example. We first need to distinguish its puzzle from the observed handedness of weak interactions.

27 Why is strong CP violation so small?

Before reading on. If the equations permit a symmetry-violating term, must its coefficient have a large nonzero value?

No. A permitted coefficient still has to be measured or explained. The strong-CP problem asks why one particular physical parameter appears to be extremely small, although ordinary QCD permits it.

The letters CP combine two transformations. Charge conjugation exchanges particles with antiparticles; parity reverses spatial handedness. Weak interactions violate CP in observed processes. That does not establish a large CP-violating effect from the strong interaction.

QCD allows a term associated with the structure of its gauge fields. Its coefficient and phases in the quark mass terms combine into a physically relevant parameter. Changing field conventions can move the phases between different parts of the description, but cannot simply remove that physical combination.

One sensitive consequence would be a permanent electric dipole moment of the neutron. A neutral object can have such a moment: zero total charge does not forbid a small displacement between positive and nega-

tive charge distributions. For a particle with spin, a permanent dipole tied to that spin probes the relevant discrete symmetries. Searches constrain it very strongly.

The puzzle is thus not that experiments have proved the coefficient exactly zero. It is that its observed bounds make it much smaller than one might have expected without a protecting mechanism. This is distinct from asking why the ordinary strong coupling has a particular strength.

The Peccei–Quinn mechanism is one proposed dynamical explanation. It introduces a new symmetry structure and an associated axion field whose dynamics can relax the problematic parameter. That is an additional physical proposal with experimental consequences, not a result of the word symmetry alone. It should not be reported as an established detection of an axion.

For a new framework, declaring all coefficients real in one preferred basis is not enough. One must show that the physically relevant phase combination is controlled by the action and remains controlled when quantum effects are included.

Try 27. Why does a tight upper bound on a neutron electric dipole moment not prove that the strong-CP parameter is exactly zero?

The strong-CP question concerned a small symmetry-violating coefficient. Another puzzle concerns whether a scalar mass can stay small when the theory contains much heavier physics.

28 Why is a light Higgs mass a puzzle?

Before reading on. Suppose two contributions, 1000 and -999, leave a total of 1. Does measuring that total explain why the two contributions nearly cancel?

The arithmetic is easy; the explanation is not supplied by it. The electroweak scale is the characteristic energy associated with the Higgs background and the W and Z masses. The hierarchy problem asks whether that scale can remain much smaller than a higher physical scale without unexplained sensitivity to the physics at that higher scale.

The Higgs is a scalar field. Its mass parameter can receive quantum contributions from fields to which it couples. If new heavy particles exist and couple to it, their threshold effects can be large compared with the measured electroweak scale. One then asks whether a symmetry or mechanism keeps the observable scalar mass small.

The example of 1000 minus 999 is only an illustration of cancellation. It is not a calculation of Higgs mass terms. Those terms depend on the model, the parameters and the renormalisation prescription, which relates calculated quantities to measured ones.

A formally divergent expression or a large cutoff is not, by itself, an observable heavy particle. The sharper physical question concerns sensitivity to an actual proposed high scale and its couplings. Assumptions about that physics therefore matter to the diagnosis.

There are proposed responses. Supersymmetry can relate bosonic and fermionic contributions when its required structure is present. Composite-Higgs models change the high-energy description. Other approaches question the demand that a small measured parameter should avoid large unexplained cancellations, often called naturalness, or propose different explanations for parameter selection. These pro-

posals have different costs and tests; none follows from the illustrative subtraction.

This question is related to, but distinct from, why gravity is weak in low-energy particle experiments. A ratio of scales describes the separation. Stability under quantum contributions asks how that separation is maintained in a specified theory.

The Standard Model can fit a Higgs mass without settling that deeper explanation. A foundation claiming to do more must identify the relevant scales and show how its mechanism controls the physical sensitivity. Replacing a continuum calculation by a finite register does not automatically do so.

Try 28. If the first contribution in the illustration changes from 1000 to 1001 while the second remains -999, what is the new total? What extra claim would a protection mechanism need to justify?

Optional

Formulas and a colour diagram

Two formulas and a diagram condense ideas we have already used. One relates a system's rest energy to its mass. The other describes the oscillation of a two-component neutrino state in an idealised setting. Neither formula will choose the physical inputs for us; each shows how specified inputs are connected.

The first question keeps track of the system and the observer's frame. The second starts with differences between ordinary numbers, then explains how the sine function describes a repeating probability. A sine is a geometrical way of reading one component of a turning arrow. It will be introduced before it appears in the oscillation formula.

A final question examines the colour triangle and distinguishes a partial charge diagram from a full quantum singlet. These pages are optional. The aim is to recognise the arithmetic inside the notation, including which quantities must be measured separately. They are not instructions to insert laboratory numbers into a formula before checking its assumptions and units.

29 How does rest energy become a mass formula?

Before reading on. A bound system releases energy to its surroundings. Should its mass be calculated from the energy it had before release or the energy it retains afterwards?

The binding example began with 100 units and released 3. The final system retained 97 units in its rest frame. That retained rest energy determines its mass. For a numerical conversion, imagine choosing a mass unit whose equivalent rest energy is 10 of our energy units. A system with 20 energy units then has mass 2 in the chosen mass units;

one with 30 has mass 3. The conversion is proportional. In ordinary units the conversion factor is the square of the vacuum light speed. The familiar formula is

$$E_{\text{rest}} = mc^2.$$

Here E_{rest} names rest energy, m names the mass of the whole system, and c is the invariant speed of light in vacuum. The factor converts mass units to energy units. It does not describe a substance called mass turning into a different substance called energy.

The word rest is essential. A moving system has kinetic energy in addition to its rest energy. Its total energy depends on the observer's frame. Its invariant mass is obtained from energy and momentum together, or directly from the energy in its centre-of-momentum frame when that frame exists.

A collection of individually massless photons can have nonzero invariant mass as a system. Two equal-energy photons moving in opposite directions have zero total momentum and positive total energy. Their combined rest energy is therefore nonzero, although neither photon has individual rest mass.

This is one reason to avoid identifying system mass with a sum of constituent rest masses. Energy in motion and fields contributes to the whole account. Binding changes that account when energy is exchanged with the surroundings.

The equation also does not calculate the binding energy by itself. To find whether the example releases 3 units or some other amount, we need the interaction theory and its states. Relativity tells us how the resulting rest energy is related to mass.

Whenever you see the formula, ask which system is being considered, which frame is used and which energy is included. Those three questions prevent more mistakes than memorising the symbols alone.

Try 29. Can two photons have a nonzero combined invariant mass while each photon has zero rest mass? Describe the example.

The mass formula referred to the energy of a whole system at rest. Neutrino oscillations use another aspect of energy: the rate at which the relative phase changes. Their formula will show why a repeating flavour pattern measures differences of squared masses.

30 Why do oscillations reveal differences of squared masses?

Before reading on. If the same number is added to every squared mass, can a measurement depending only on their differences detect that addition?

Take the illustrative lists 1, 4, 10 and 6, 9, 15. In either list the successive differences are 3 and 6, and the widest difference is 9. A common offset changes the entries without changing those differences. Here squared mass means each mass multiplied by itself; a difference of squared masses is not the square of a mass difference.

The relativistic energy relation explains why squared masses appear. For a very fast particle, the leading energy contribution is set by its momentum. The leading small correction due to mass is proportional to mass squared divided by twice that energy. Comparing two propagating components therefore gives an energy or phase difference involving the difference of their squared masses. The coherent production-to-detection calculation gives the same leading oscillation phase without choosing a measurable separate flight time for each mass component.

Now recall the equal-mixture example. Conversion probabilities were 0, one half, 1, one half and 0 as the relative phase went through zero, quarter, half, three-quarter and full turns. A sine function is a compact way to describe that repeating shape. For a unit-length arrow making an angle with the horizontal axis, its vertical component is the sine of that angle. Squaring this component gives a number between zero and one.

For flavour conversion we use the sine of half the relative phase angle, then square it. A relative half-turn makes that half-angle a quarter-turn, whose sine is 1. A relative full turn makes the half-angle a half-turn,

whose sine is zero. The half-angle is why one complete relative-phase turn gives one complete probability cycle.




Figure 8: The ideal equal-mixture, two-flavour conversion probability. The horizontal axis is relative phase, not distance in metres. The formula relates that phase to distance, energy and a mass-squared difference.

An unequal mixture reduces the largest possible conversion. Its mixing angle is written θ . For the equal mixture it is 45 degrees. Twice that angle is 90 degrees, giving a squared sine of 1. The leading vacuum formula is

$$P_{\text{change}} = \sin^2(2\theta) \sin^2\left(\frac{\Delta m^2 L}{4E}\right).$$

The first factor sets the maximum conversion. The second supplies the repeating phase dependence. The notation sin squared means calculate the sine and then square the answer. It does not mean take the sine of a squared angle.

The symbol Δm^2 names a difference of squared masses, L is the travel distance and E the leading neutrino energy. In this formula the speed of light and reduced Planck constant have been set to one by a consistent choice of units. Mass and energy then use the same unit, and distance uses the reciprocal energy unit. The angle in the second sine is in radians, where a half-turn is pi radians. The fraction in that sine is half the relative propagation phase.

These conventions are not permission to put a distance in metres and an energy in joules straight into the printed fraction. The physical constants or equivalent conversion factors must be restored for those units. Here the formula is used to inspect the dependence: doubling both distance and energy leaves the phase unchanged.

The expression assumes coherent, highly relativistic propagation in a two-flavour vacuum model, keeping the leading mass-dependent terms. Real analyses use three-flavour mixing, energy distributions, matter effects and coherence conditions. Within the stated approximation, a common squared-mass offset is invisible to the oscillation probability. Other measurements are needed for the absolute mass scale.

Try 30. Within the stated approximation, what happens to the oscillation phase if distance and energy are both doubled?

Complex amplitudes and internal colour labels are different kinds of information. A tempting picture gives the three colours three arrows 120 degrees apart. We can now ask what that picture captures and what a true colour-neutral quantum state requires.

31 Can three colour arrows explain a proton?

Before reading on. Three equal arrows 120 degrees apart sum to zero. Does that arithmetic prove their quantum state is unchanged by every colour transformation?

It does not. The arrow cancellation is correct geometry, but it contains less information than the full colour transformation law. Confusing the two would turn a useful illustration into a false derivation of the strong interaction.

A quark's colour state has three complex amplitudes relative to three chosen basis labels: red, green and blue. These basis states are mutually distinguishable within the internal-state mathematics. They are not three phases of one complex number. A colour transformation can mix the three amplitudes.

There is nevertheless a legitimate triangle picture. Choose two commuting colour generators. A generator describes one small change of colour state; two commuting generators are compatible, so their charge values can be specified together in a chosen basis. Their values for the three basis colours form a triangle in a two-dimensional diagram called a weight diagram. Its directions represent those charge values, not phases on our earlier interferometer calculation page.




Figure 9: A schematic colour weight diagram. The three positions sum to the centre. This records two compatible colour charges; it does not display all eight colour generators or test invariance under all of them.

Requiring the two displayed charges to sum to zero is weaker than requiring a singlet, a state unchanged by every colour transformation. The eight-dimensional gluon colour description is called an octet. Two of its directions have zero values for both charges shown here, yet they are not singlets: other colour transformations mix them with the other octet components.

For a meson’s simple quark–antiquark colour singlet, the state is an equal coherent combination of red with anti-red, green with anti-green and blue with anti-blue. A single red–anti-red product is not by itself that invariant state. For three quarks, the singlet combines the six permutations of red, green and blue with alternating signs and proper normalisation.

A determinant is an alternating sum of products of entries in a square matrix. Its use for the three-quark combination records the sign change when two quarks are exchanged. It is not a claim that three ordinary waves simply cancel. Likewise, an antisymmetric two-quark colour combination transforms as an antitriplet, a three-component colour description obeying the same transformation rule as an antiquark; an arbitrary green–blue product does not automatically behave as a single anti-red state.

These facts explain why the proposed triangle is a useful shadow of the algebra, but not its replacement. They also leave the dynamics to be supplied: an invariant colour combination need not be a stable bound particle.

Try 31. Does testing the two charges in the triangle suffice to establish a colour singlet? What stronger condition is needed?

For checking

Answers to the Try questions

Try the questions before reading these answers. Compare the physical assumptions and the first step of the reasoning, not just the final number. A matching number can hide a different interpretation.

1. The surroundings received the released energy. The bound system retains 97 units; the combined system-plus-surroundings account still balances.
2. No. The quoted constraint assumes a mass range and a coupling to the Z. A different particle requires the appropriate separate test.
3. The three-component quark representation and SU(3) colour symmetry are supplied. The eight gluon generators then follow.
4. The proton total is +1. The neutron total is 0. These are valence charge accounts, not claims that the internal states are static beads.
5. The difference is 0.036. A visually small difference can still be much larger than the experimental uncertainty.
6. It doubles from 0.3 to 0.6 newtons, within the stated fixed-field, perpendicular-wire idealisation.
7. The available channel needs a transition amplitude. Couplings, selection rules and phase space determine its rate.
8. The sum is +1. Energy, momentum, angular momentum or another relevant conserved quantity must also be checked.
9. Two equal opposite photon momenta cancel. One photon has nonzero momentum, so it cannot match a positive-energy initial system with zero total momentum by itself.
10. Helicity. Overtaking a massive particle can reverse its momentum direction relative to the observer while leaving the relevant spin orientation unchanged.
11. $1 - 1 + 0 = 0$. The proton's positive charge is balanced by the electron's negative charge.

12. An internal field contribution. The effective low-energy interaction does not require a real W of reduced mass.
13. The natural width doubles, assuming the standard lifetime–width relation. An observed detector width can have other contributions.
14. The coefficient becomes one quarter as large. Its squared contribution becomes one sixteenth, with the other factors held fixed.
15. Twenty events on average. Individual runs fluctuate, and real detectors have efficiencies and backgrounds.
16. The other flavour has probability 1 in the ideal equal-mixture model. The component masses remain unchanged.
17. Both lists have successive differences 3 and 6. A common offset in squared mass is invisible to these differences.
18. The emitted electron contributes -1, so $18 - 1 = 17$. A precise energy calculation must also use consistent nuclear or atomic masses.
19. The 12-unit final state is too energetic. The 8-unit final state has an incompatible total charge. Lower energy alone does not make it accessible.
20. Allowed channels, coupling amplitudes and available phase space. The initial and final states must be specified too.
21. Two units. The transition must still satisfy the other selection rules, and its matrix element is needed to assign a rate.
22. The nucleon count changes from 90 to 92; electric charge remains +40. The count alone does not determine the new isotope's stability.
23. Two units remain. Energy availability does not supply the interaction amplitude, final-state conditions or rate.
24. The observable charge and energy balance must agree. Diagram arrows encode the representation and charge-flow convention.
25. Twelve: three polarisations for each of the two W states and the Z give nine, the photon contributes two, and the Higgs contributes one. This is a degree-of-freedom count, not a dynamical derivation.
26. There are twelve gauge fields in that local count. Their number does not fix the coupling coefficients.

- 27.** An upper bound leaves sufficiently small nonzero values possible. It is not a measurement of exact zero.
- 28.** The total becomes 2. A protection mechanism must explain the small physical result and its stability under the changes permitted by the actual theory; the toy subtraction does not do that.
- 29.** Yes. Two equal-energy, oppositely moving photons have positive total energy and zero total momentum, so the combined system has nonzero invariant mass.
- 30.** It is unchanged, since the distance-to-energy ratio is unchanged.
- 31.** No. A colour singlet must be invariant under the full colour group, not just have zero values of two selected commuting charges.

For returning to

Words and further reading

Words to return to

Coupling

A parameter in the theory that sets the strength of a specified interaction.

Flavour

The electron, muon or tau alternative selected by weak neutrino production and detection. It is distinct from a definite-mass propagation state.

Colour singlet

A complete strong-interaction state unchanged by colour transformations. This is not a mixture of visible paint colours.

On shell

Satisfying the energy-momentum relation of a physical free-particle state. An internal line in a calculation is not a detected free particle.

Phase space

The set of final momenta and states permitted by the conservation conditions; its available extent enters reaction rates.

Selection rule

A restriction on a transition arising from symmetries, conserved quantities or the states involved.

Vacuum

A reference state of the fields, usually a lowest-energy state in the specified setting; it need not be devoid of correlations.

What to consult next

For experimental particle properties, the Particle Data Group provides a catalogue and topical reviews. Its neutrino review separates oscillation parameters from absolute masses and the Dirac–Majorana question.[1, 2]

CERN’s Higgs introduction is a route into the electroweak mass mechanism. It should be read alongside the distinction made here between elementary mass terms and the mass of composite hadrons.[3] For further

mathematical work, the quantum notes cited in the first volume provide the state and amplitude tools needed before a field-theory course.

This volume deliberately uses rounded comparisons and invented arithmetic examples rather than precision parameter fits. Its numbers do not constitute new measurements or framework predictions. When consulting precision bounds, check the dataset, confidence level and model assumptions beside each bound.

The field-theory discussion uses Tong’s particle-physics notes and the PDG’s QCD and Higgs reviews.[4, 5, 6] They distinguish the tested formulation from its unsettled parameters.

References

- [1] Particle Data Group, *Review of Particle Physics*, 2025 update. [Review homepage](#); [topical reviews](#). This fixed edition is used rather than silently mixing later online tables.
- [2] M. C. Gonzalez-Garcia and R. Wendell, Neutrino Masses, Mixing, and Oscillations, in the 2025 *Review of Particle Physics*. [Review PDF](#).
- [3] CERN, The Higgs boson. [CERN physics introduction](#).
- [4] D. Tong, *Particle Physics*, lecture notes, especially the Standard Model and What We Don’t Know. [Author’s notes](#).
- [5] Particle Data Group, Quantum Chromodynamics, in the 2025 *Review of Particle Physics*. [Review PDF](#).
- [6] Particle Data Group, Status of Higgs Boson Physics, in the 2025 *Review of Particle Physics*. [Review PDF](#).