AGH University of Krakow

Faculty of Physics and Applied Computer Science

MODULE 6

ELEMENTS OF STANDARD MODEL PHENOMENOLOGY

Supporting Lecture Notes for *The Standard Model*

The Guideline

The aim of this expanded module note is to show how the formal structure of the Standard Model becomes experimentally testable. Starting from interaction terms in the Lagrangian, the note develops the logic of Feynman rules, Feynman diagrams, amplitudes, squared amplitudes, phase space, decay widths, branching ratios, scattering cross sections, and selected observables that are central in particle-physics experiments. Throughout, the emphasis is on interpretation rather than excessive technical detail: the purpose is not to turn the module into a full quantum field theory or collider-analysis course, but to make clear how predictions are extracted from the theory and why measurable rates and distributions are the natural bridge between the Standard Model as a gauge theory and the world of experimental data.

Prepared for the Standard Model course

Hamzeh Khanpour

`hamzeh.khanpour@fis.agh.edu.pl`

June 8, 2026

Contents

Preface

These supporting notes are written for *Module 6: Elements of Standard Model Phenomenology* of the AGH course *The Standard Model*. The central pedagogical question of the module is: how does one pass from the formal language of fields, symmetries, and interaction terms to quantities that can actually be measured? Earlier modules established the relativistic and gauge-theoretic structure of the Standard Model, developed the strong-interaction sector, and explained electroweak symmetry breaking and Higgs couplings. Module 6 now takes the next essential step. It is the first module whose primary task is not to build the theory, but to show how the theory becomes predictive in terms of decays, scattering rates, and physically useful observables.

The pedagogical logic of the module is

$$\begin{aligned}\text{interaction structure} &\longrightarrow \text{Feynman rules and diagrams} \longrightarrow \text{amplitudes} \longrightarrow |\mathcal{M}|^2 \\ &\longrightarrow \text{phase space} \longrightarrow \Gamma \text{ and } \sigma \longrightarrow \text{branching ratios and distributions} \longrightarrow \text{experiment.}\end{aligned}$$

This chain is simple to write, but each arrow carries important conceptual content. The Lagrangian does not directly hand us event counts. Instead, it determines vertices and propagators, which determine amplitudes, which in turn determine rates once kinematics and phase space are taken into account. The experimental language of decays, lifetimes, cross sections, invariant-mass peaks, angular distributions, and event yields is therefore not an external add-on to the Standard Model. It is how the theory becomes falsifiable.

The note is intentionally placed after Modules 1–5. It assumes the student has already seen Lorentz symmetry, spinors, gauge invariance, the Standard Model gauge structure, QCD, and the Higgs mechanism. For that reason the present note does not attempt to rederive quantum field theory from first principles, nor does it attempt to become a full amplitudes course, a full renormalisation chapter, or a full collider-analysis manual. Its scope is deliberately bounded. The emphasis is on tree-level structure, physically transparent examples, and the conceptual route from formal expressions to measurable quantities.

The note has three closely connected aims. The first is structural: to explain what Feynman rules and Feynman diagrams encode and how amplitudes are built. The second is phenomenological: to introduce decay widths, branching ratios, cross sections, and selected differential observables as the natural language of Standard Model predictions. The third is bridge-like: to prepare the student for later modules in which open questions, data-analysis workflows, and laboratory interpretation become central. Module 6 should therefore be read as the place where the Standard Model begins to look not only mathematically constrained, but experimentally operational.

Conventions and notation

- Metric signature: $\eta_{\mu\nu} = \mathrm{diag}(1, -1, -1, -1)$.
- Natural units: $\hbar = c = 1$.
- Four-momentum notation: a particle momentum is written $p^\mu = (E, \mathbf{p})$, with $p^2 = m^2$ for an on-shell particle.

- Mandelstam variables for $2 \to 2$ scattering:

$$s = (p_1 + p_2)^2, \qquad t = (p_1 - p_3)^2, \qquad u = (p_1 - p_4)^2.$$

- Scattering or decay amplitudes are denoted by $\mathcal{M}$ or, when convenient, by $\mathcal{A}$.
- The quantity relevant for unpolarised rates is typically a spin-summed or spin-averaged squared amplitude, schematically written as $\overline{|\mathcal{M}|^2}$.
- Lorentz-invariant phase space for a final state with particles i is written schematically as

$$\mathrm{d}\Phi_f = (2\pi)^4 \delta^{(4)}\left(p_{\text{in}} - \sum_i p_i\right) \prod_i \frac{\mathrm{d}^3 p_i}{(2\pi)^3 2E_i}.$$

- The total decay width of a particle A is denoted Γ_A and is related to the lifetime by

$$\tau_A = \frac{1}{\Gamma_A}.$$

- A partial width into channel i is written Γ_i, and the branching ratio is

$$\mathrm{BR}_i = \frac{\Gamma_i}{\Gamma_{\text{tot}}}.$$

- Total cross sections are denoted by σ and differential cross sections by expressions such as $\mathrm{d}\sigma/\mathrm{d}\Omega$, $\mathrm{d}\sigma/\mathrm{d}t$, or $\mathrm{d}\sigma/\mathrm{d}m$.
- When hadron-collider observables are mentioned at orientation level, p_T denotes transverse momentum, η pseudorapidity, and m_{ab} the invariant mass of an identified subsystem.

1 Introduction and module roadmap

1.1 Why phenomenology deserves its own module

Phenomenology deserves its own module because a quantum field theory is not tested directly by staring at its Lagrangian. What experiments actually measure are event rates, branching fractions, invariant-mass spectra, angular distributions, and counting observables built from reconstructed particles. The Standard Model would not be a predictive physical theory if it did not contain a disciplined route from formal interaction terms to these measurable quantities. Module 6 is devoted to making that route visible.

This is why the module should not be interpreted merely as a collection of formulas for Γ and σ. Its deeper role is to explain how the Standard Model becomes experimentally meaningful. Feynman diagrams, amplitudes, decay widths, and cross sections are not disconnected technical tricks. They are the language that connects the theory to collider experiments, fixed-target experiments, and precision measurements.

1.2 From Module 3, Module 4, and Module 5 to Module 6

Module 3 developed the gauge-theory logic of the Standard Model and identified the allowed interaction structure. Module 4 then showed how one part of that structure — QCD — leads to a rich phenomenology of partons, jets, and strong-interaction processes. Module 5 explained how electroweak symmetry breaking and Yukawa interactions determine masses and couplings. Module 6 now asks what one can actually do with these ingredients once they are known.

In this sense the module sits exactly where it should in the course architecture. The earlier modules established the kinematic language, the matter content, the gauge structure, and the Higgs sector. The present module turns these into processes. The couplings derived earlier now become entries in vertices, propagators, and amplitudes. The result is the first systematic discussion of how Standard Model predictions are turned into quantities that can be compared with data.

1.3 What this module will and will not do

This module will introduce the interpretation of Feynman rules and Feynman diagrams, the role of amplitudes, the transition from amplitudes to squared amplitudes, the importance of phase space, the meaning of decay widths and branching ratios, the logic of scattering cross sections, and the physical content of selected observables used in particle-physics experiments. It will also use a few canonical tree-level examples to make the formalism concrete.

At the same time, the module has deliberately bounded scope. It will not present a full LSZ derivation of the S-matrix, a full loop-calculation programme, a full renormalisation chapter, a detailed treatment of trace technology, or a full detector and statistics course. Those belong either to later modules, to advanced quantum field theory, or to laboratory-oriented study. The present note stays at the level where the main logic is visible and physically interpretable.

1.4 Roadmap of the note

The note begins by explaining why phenomenology is a necessary bridge in the Standard Model and by identifying the role of perturbative reasoning at tree level. It then turns to interaction terms, Feynman rules, and Feynman diagrams, stressing both what these objects encode and what they do not literally mean. The next step is the amplitude itself and the reason why physical rates depend on $|\mathcal{M}|^2$ rather than on $\mathcal{M}$ directly. Kinematics, phase space, decays, branching ratios, and cross sections are then introduced in a progressive way.

The later sections connect this structure to a few canonical processes, to resonances and intermediate propagators, and finally to collider-facing observables such as invariant masses and transverse-momentum distributions. The concluding sections make explicit how Module 6 prepares the transition to Standard Model limitations and to data-analysis modules.

Compact roadmap. The logical flow of the module is

$$\begin{aligned}\text{interaction terms} &\longrightarrow \text{Feynman rules and diagrams} \longrightarrow \mathcal{M} \longrightarrow \overline{|\mathcal{M}|^2}\\ &\longrightarrow \text{phase space} \longrightarrow \Gamma \text{ and } \sigma\\ &\longrightarrow \text{branching ratios and observables} \longrightarrow \text{experimental interpretation.}\end{aligned}$$

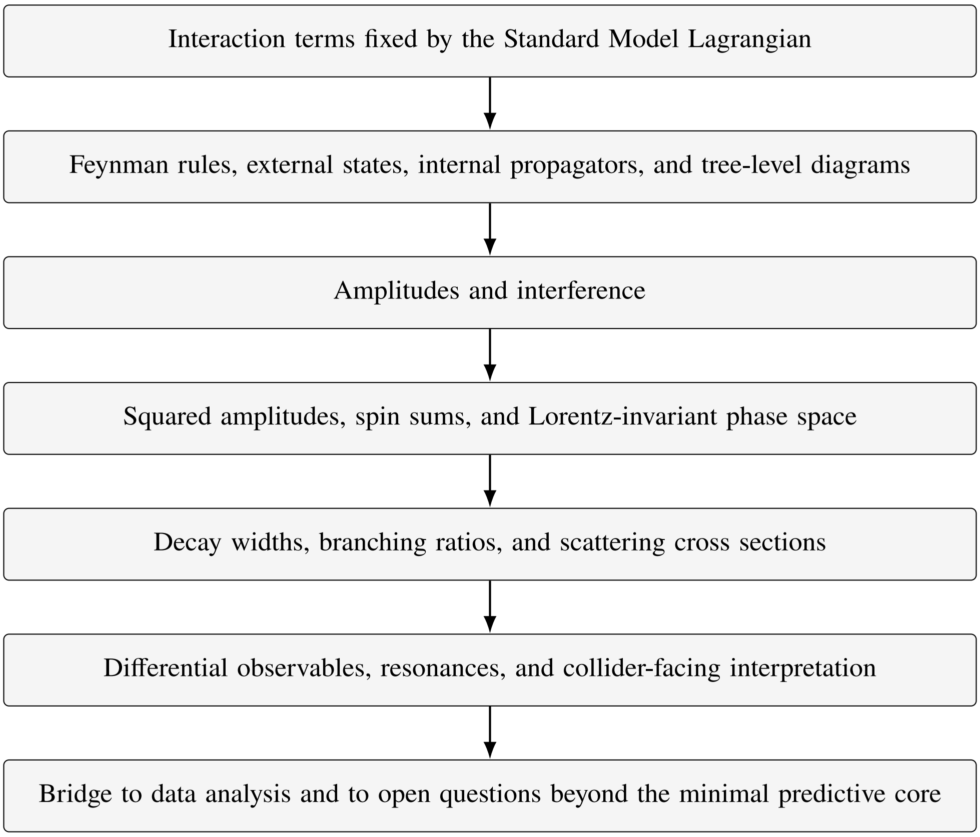

Figure 1: Conceptual roadmap of Module 6. The note explains how one moves from interaction terms in the Standard Model to amplitudes, rates, differential observables, and their experimental interpretation.

2 Why phenomenology is a necessary bridge in the Standard Model

2.1 From Lagrangians to measurable quantities

A Lagrangian is the compact encoding of the dynamical content of a field theory. It tells us which fields exist, which symmetries are respected, and which interactions are allowed. Yet a detector does not measure a Lagrangian density. It records hits, tracks, calorimetric deposits, reconstructed particles, and event counts in certain channels. The bridge between these worlds is built by amplitudes and observables.

For this reason, phenomenology should not be regarded as something secondary to the theory. It is the part of the theory that knows how to talk to experiment. A Standard Model interaction term matters physically because it contributes to a process, and that process matters because it changes a rate or a distribution that can be tested.

2.2 Why experiments measure rates, distributions, and event counts rather than "the Lagrangian"

Experimentally one asks questions of the following kind: how often does a particle decay into a given final state; what is the probability that two incoming particles scatter into a specified channel; where does the invariant mass of the final state cluster; how is the angular distribution shaped; and how many events are expected after a given selection? These are questions about observables.

The formal theory must therefore be translated into objects that answer such questions. This translation does not erase the underlying theory; rather, it organises it. One takes the interaction structure, constructs amplitudes, squares and sums appropriately, integrates over the allowed kinematics, and finally obtains decay widths, branching ratios, or cross sections. In this sense rates and distributions are the operational face of the Lagrangian.

2.3 The central role of amplitudes, decays, and cross sections

Among the many possible observables in particle physics, decay widths and scattering cross sections occupy a privileged place. A decay width measures the instability of a particle and therefore the strength and kinematic accessibility of its available channels. A cross section measures the effective rate at which an interaction occurs in a given initial state. Both are derived from the same central object: the quantum-mechanical amplitude.

It is therefore useful to think of Module 6 as organised around one question: given an interaction term, how does one obtain a prediction for a decay width or a cross section? The answer involves several intermediate ideas — diagrams, propagators, interference, phase space, and flux factors — but the logic remains unified.

2.4 A short methodological remark on perturbation theory and tree level

The present module works predominantly at tree level. This is not because higher-order corrections are unimportant in modern particle physics, but because tree-level reasoning is where the conceptual anatomy of a process is easiest to see. At tree level one can identify which coupling enters, which propagator appears, what the relevant kinematic variables are, and why one channel may dominate over another.

Higher-order corrections improve precision, modify normalisations, and sometimes change shapes significantly. But pedagogically they come after the basic structure is understood. For a first phenomenology bridge, tree-level processes are the correct default language.

Take-home message

Phenomenology is the step in which the Standard Model stops looking like a catalogue of fields and starts looking like a predictive theory. The Lagrangian determines interactions, but amplitudes, widths, cross sections, and distributions are the quantities through which those interactions become experimentally testable.

3 From interaction terms to Feynman rules

3.1 Interaction terms as the source of vertices

The starting point for perturbative phenomenology is the decomposition of the Lagrangian into free and interaction pieces,

$$\mathcal{L} = \mathcal{L}_{\text{free}} + \mathcal{L}_{\text{int}}. \tag{3.1}$$

The free part determines which fields propagate and therefore fixes the propagators and the on-shell interpretation of external particles. The interaction part determines which fields can meet at a vertex and therefore fixes the possible processes. This is the first genuinely phenomenological lesson: once the field content and the allowed interaction terms are known, one already knows a great deal about what can and cannot happen experimentally.

In the Standard Model, this separation is conceptually straightforward even when the full Lagrangian is lengthy. One may write schematically

$$\mathcal{L}_{\mathrm{SM}} = \mathcal{L}_{\text{gauge kin}} + \mathcal{L}_{\text{fermion kin}} + \mathcal{L}_{\text{Higgs}} + \mathcal{L}_{\text{Yukawa}}. \tag{3.2}$$

After expanding covariant derivatives and, in the non-Abelian case, the field strengths, all terms quadratic in the fields belong to the free theory while terms cubic or quartic in the fields describe interactions. The passage from the earlier structural modules to the present phenomenology module is precisely the passage from this decomposition to observable processes.

3.2 The Standard Model Lagrangian as the source of phenomenology

The previous modules developed the gauge principle, the electroweak structure, QCD, and the Higgs sector. For phenomenology, one does not need to rewrite all of that from scratch, but one should keep visible which pieces of the Standard Model Lagrangian feed directly into measurable processes. A useful selective summary is

$$\begin{aligned}\mathcal{L}_{\mathrm{SM}} \supset & -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}Z_{\mu\nu}Z^{\mu\nu} - \frac{1}{2}W^{+}_{\mu\nu}W^{-\,\mu\nu} - \frac{1}{4}G^{a}_{\mu\nu}G^{a\mu\nu} \\ & + \sum_f \bar{f}(\mathrm{i}\gamma^\mu\partial_\mu - m_f)f + \mathcal{L}^{\gamma}_{\mathrm{int}} + \mathcal{L}^{W}_{\mathrm{int}} + \mathcal{L}^{Z}_{\mathrm{int}} + \mathcal{L}^{g}_{\mathrm{int}} + \mathcal{L}^{h}_{\mathrm{int}} + \cdots .\end{aligned} \tag{3.3}$$

Here the ellipsis suppresses terms that are structurally important but not needed every time one discusses a given process. The key point is that phenomenology does not start from a diagrammatic picture floating free in space. It starts from specific interaction terms already determined by the theory.

For later use, it is helpful to display a few representative couplings after electroweak symmetry breaking:

$$\mathcal{L}^{\gamma}_{\mathrm{int}} = -e\sum_f Q_f\, \bar{f}\gamma^\mu f\, A_\mu, \tag{3.4}$$

$$\mathcal{L}^{W}_{\mathrm{int}} = -\frac{g}{\sqrt{2}}\left(W^+_\mu\, \bar{u}_i\gamma^\mu\frac{1-\gamma^5}{2}V_{ij}d_j + W^+_\mu\, \bar{\nu}_\ell\gamma^\mu\frac{1-\gamma^5}{2}\ell + \mathrm{h.c.}\right), \tag{3.5}$$

$$\mathcal{L}^{Z}_{\mathrm{int}} = -\frac{g}{2\cos\theta_W}Z_\mu\sum_f \bar{f}\gamma^\mu\left(g^f_V - g^f_A\gamma^5\right)f, \tag{3.6}$$

$$\mathcal{L}^{h}_{\mathrm{int}} = -\sum_f \frac{m_f}{v}\, h\,\bar{f}f + \cdots . \tag{3.7}$$

These are precisely the kinds of terms that later appear as vertices in scattering and decay amplitudes. Their coefficients tell us which couplings are involved, which currents are chiral, which flavour-changing structures are allowed, and how masses enter.

Take-home message

The Standard Model Lagrangian is already a phenomenology generator. Once the free and interaction parts are identified, the route to Feynman rules, amplitudes, widths, and cross sections is conceptually fixed.

3.3 External lines, internal lines, and vertices

At introductory level it is useful to distinguish three kinds of ingredients in a Feynman diagram. External lines correspond to asymptotic particles in the initial or final state. Internal lines correspond to propagators of fields that are exchanged virtually. Vertices encode the coupling structure extracted from the interaction Lagrangian.

This distinction is conceptually important because only external particles are interpreted as observed or prepared states. Internal lines are not directly observed states; they belong to the intermediate quantum-mechanical structure of the perturbative contribution. This point becomes especially important when students first encounter the language of "virtual particles".

3.4 What a Feynman rule encodes

A Feynman rule is not merely a picture-drawing convention. It is an algebraic instruction associated with a specific element of the diagram. Vertices contribute coupling factors and Lorentz or spinor structure. Internal lines contribute propagator factors. External lines contribute wavefunctions or polarisation vectors, depending on the particle species. The full amplitude for a given diagram is obtained by multiplying these elements together and contracting the relevant indices.

Definition 3.1: Feynman rule

A Feynman rule is the algebraic factor assigned to a propagator, vertex, or external state in a perturbative contribution to a process. Together, the full set of rules allows one to translate a diagram into a corresponding mathematical expression for the amplitude.

With a common set of conventions, some standard examples are

$$\text{fermion propagator:} \quad \frac{\mathrm{i}(\gamma^\mu q_\mu + m)}{q^2 - m^2 + \mathrm{i}\epsilon}, \tag{3.8}$$

$$\text{photon propagator (Feynman gauge):} \quad \frac{-\mathrm{i}g_{\mu\nu}}{q^2 + \mathrm{i}\epsilon}, \tag{3.9}$$

$$\text{QED vertex:} \quad -\mathrm{i}eQ_f\gamma^\mu. \tag{3.10}$$

The precise signs and factors of i depend on convention, but the structural message is always the same: the interaction term fixes the vertex factor, and the free part fixes the propagator.

3.5 QED as the simplest prototype

Quantum electrodynamics is the cleanest first example because the interaction structure is simple and familiar. The QED interaction term is

$$\mathcal{L}_{\text{int}}^{\text{QED}} = -e\,\bar{\psi}\gamma^{\mu}\psi A_{\mu}. \tag{3.11}$$

This implies a fermion–fermion–photon vertex. At tree level, many standard processes can be understood from repeated use of this one structure. This makes QED the natural pedagogical entry point even when one ultimately wants to discuss weak decays or Standard Model processes more broadly.

Example 3.1: A vertex read directly from the interaction term

The QED interaction term contains one photon field and two fermion fields. This already tells us the topology of the elementary interaction: a charged fermion line can emit or absorb a photon. In a standard convention the corresponding vertex factor is $-\mathrm{i}e\gamma^{\mu}$ for a unit-charged fermion. The diagrammatic picture and the algebraic rule therefore come from the same source.

3.6 Electroweak and QCD remarks at orientation level

The same logic extends beyond QED. In the electroweak theory, chiral couplings and massive weak bosons lead to a richer vertex structure. In QCD, colour indices and non-Abelian self-interactions enlarge the range of possible vertices. But the conceptual grammar does not change: the interaction terms determine which vertices can appear, and the propagators determine how internal fields are exchanged.

In particular, the weak charged current is phenomenologically special because it changes flavour. The term $W_{\mu}^{+}\,\bar{\nu}_{\ell}\gamma^{\mu}(1-\gamma^{5})\ell$ opens decay channels such as muon decay, while the quark charged current brings in CKM mixing. Likewise, the non-Abelian structure of QCD means that gluons couple not only to quarks but also to one another. Later collider phenomenology is rich partly because the underlying Lagrangian already contains these possibilities.

At the level of Module 6, the purpose is not to list every Standard Model rule exhaustively. It is to make clear that the diagrammatic and algebraic structure is not arbitrary. It is inherited from the Lagrangian and therefore from the symmetry logic developed in the earlier modules.

4 Feynman diagrams: meaning, use, and limitations

4.1 Diagrams as bookkeeping devices in perturbation theory

Feynman diagrams provide a visual and calculational language for perturbation theory. They tell us which contributions appear at a given order in the coupling expansion and help organise the algebra of the amplitude. For this reason they are powerful both conceptually and practically.

Yet the key word is *organise*. A diagram is not the full process; it is one perturbative contribution to the full amplitude. In many important introductory examples, the lowest-order contribution dominates strongly enough that a single diagram captures most of the physics. But in principle the physical amplitude is a sum over all relevant contributions of the specified order, and beyond that over higher orders as well.

4.2 Tree-level diagrams versus higher-order corrections

Tree-level diagrams are those without closed loops. They are the simplest contributions and typically the first ones studied in a course because they already exhibit the core dependence on couplings, propagators, and kinematics. Higher-order diagrams contain loops and bring in corrections that improve the precision of the prediction.

At tree level, one often sees clearly why a particular channel exists, which mediator is involved, and how the order in the coupling counts. For example, a QED process with two electromagnetic vertices carries a lowest-order amplitude proportional to e^2 and a rate proportional to $e^4 \sim \alpha^2$. This simple counting is one reason why tree-level diagrams are pedagogically valuable.

4.3 What diagrams do and do not literally represent

Students often encounter a useful but dangerous phrase: a diagram "shows what happens. In one limited sense this is true. The diagram shows which external states, vertices, and mediators participate in a perturbative contribution. But it should not be interpreted as a classical spacetime movie in which particles follow sharp trajectories and literally toss mediators back and forth.

Remark 4.1: Do not over-literalise the picture

A Feynman diagram is best understood as a compact encoding of a term in the perturbative expansion of the amplitude. It is not a classical drawing of microscopic motion. Treating it too literally obscures interference, quantum superposition, and the fact that different diagrams contributing to the same process must be added at the amplitude level.

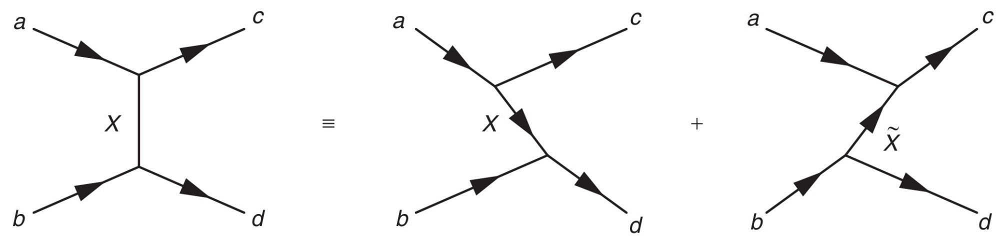

Figure 2: A Feynman diagram should be understood as a compact representation of a perturbative contribution to the amplitude, not as a literal space-time movie. The diagram for $a + b \to c + d$ shown on the left corresponds to the sum of the two possible time-ordered contributions shown on the right. This is why the physical process is obtained by adding amplitudes associated with the relevant contributions before squaring.

4.4 Virtual particles as internal lines: course-level interpretation

Internal lines are often described as virtual particles. This language is acceptable at course level provided one remembers its limitations. The internal propagator carries momentum determined by the kinematics of the diagram, but the corresponding intermediate state need not satisfy the on-shell relation $q^2 = m^2$. That is why it is not interpreted as an observed asymptotic particle.

The pedagogical value of the term *virtual* is that it distinguishes internal propagating structure from real detected particles. Its danger is that it may suggest that one is temporarily producing ordinary particles and then hiding them again. A more careful statement is that internal lines represent the propagator factors associated with fields exchanged in the perturbative contribution.

4.5 Why the simplest diagrams often dominate at introductory level

In weakly coupled theories the contributions with the fewest vertices are often numerically the most important. This is the basis for the common introductory practice of analysing the simplest tree-level diagram first. The dominance is not absolute in all circumstances, but it is sufficiently robust in many classic examples that it forms the right pedagogical starting point.

This point also helps explain why Feynman diagrams are so effective in first contact with phenomenology. One can often identify the qualitative behaviour of a decay or scattering process simply from the simplest allowed diagram, even before precision refinements are considered.

Guided checks

- Can you explain why a Feynman diagram is a contribution to an amplitude rather than the whole physical process?
- Can you state clearly the difference between an external line and an internal line?
- Can you explain why “virtual particle is a useful but limited phrase?

5 The scattering or decay amplitude as the central object

5.1 What the invariant amplitude means

The invariant amplitude $\mathcal{M}$ is the central quantum-mechanical object from which rates are derived. It contains the dynamical information associated with the chosen process: the couplings, the propagators, the spinor or polarisation structure, and the dependence on kinematic invariants. In most practical situations the amplitude is first computed for a given diagram and then summed with the amplitudes of all other contributing diagrams.

A useful schematic way to write the relation between the S-matrix and the invariant amplitude is

$$S_{fi} = \delta_{fi} + \mathrm{i}(2\pi)^4\delta^{(4)}(p_f - p_i)\,\mathcal{M}_{fi}. \tag{5.1}$$

At the level of the present module, the main message is that $\mathcal{M}$ is the quantity that carries the actual process information. Once it is known, the rate follows after squaring, summing or averaging over unobserved quantum numbers, and integrating over phase space.

5.2 Why amplitudes add before probabilities

Quantum mechanics requires that one add amplitudes, not probabilities, when different indistinguishable contributions lead to the same final state. If two diagrams contribute to the same process, the total amplitude is

$$\mathcal{M}_{\text{tot}} = \mathcal{M}_1 + \mathcal{M}_2 + \cdots . \tag{5.2}$$

The corresponding probability-like quantity is therefore

$$|\mathcal{M}_{\text{tot}}|^2 = |\mathcal{M}_1|^2 + |\mathcal{M}_2|^2 + 2\,\text{Re}\big(\mathcal{M}_1\mathcal{M}_2^*\big) + \cdots . \tag{5.3}$$

The interference term is a genuinely quantum effect. It can enhance or suppress a rate and is one of the clearest reasons why a Feynman diagram cannot be treated as a classical micro-history.

5.3 From interference to physical rates

Interference is not an optional refinement. In many important Standard Model processes it is part of the core physics. Photon and Z exchange can interfere in e^+e^- annihilation. Different helicity structures can interfere. Signal and background can interfere near a resonance. The basic rule is always the same: whenever the final state is the same and the amplitudes are coherent, they must be added before squaring.

At introductory level it is often enough to remember the algebraic pattern

$$|\mathcal{M}_1 + \mathcal{M}_2|^2 = |\mathcal{M}_1|^2 + |\mathcal{M}_2|^2 + 2\Re(\mathcal{M}_1\mathcal{M}_2^*), \tag{5.4}$$

and to interpret the cross term physically. This already teaches why apparently small additional contributions can matter.

5.4 The schematic chain $\mathcal{M} \to |\mathcal{M}|^2 \to$ observable quantity

The amplitude itself is not directly observed. What enters widths and cross sections is $|\mathcal{M}|^2$, more precisely a spin-summed or spin-averaged version of it. The broad logic is

$$\mathcal{M} \longrightarrow \overline{|\mathcal{M}|^2} \longrightarrow \text{rate density} \longrightarrow \text{integrated or differential observable.} \tag{5.5}$$

The intermediate stages matter. One must include the correct sums or averages over quantum numbers, the correct phase-space measure, and the correct normalisation of the initial or decaying state.

Take-home message

The amplitude is the dynamical heart of the process. Everything else — partial widths, total rates, branching ratios, and differential distributions — is built by turning that amplitude into a physically normalised observable.

6 Squared amplitudes, spin sums, and averages at introductory level

6.1 Why experiments are sensitive to $|\mathcal{M}|^2$

The direct output of a perturbative calculation is the amplitude, but the quantity that controls rates is its modulus squared. This reflects a general quantum-mechanical principle: probabilities are associated with squared amplitudes. In field theory the result must still be combined with the relevant kinematic measure, but $|\mathcal{M}|^2$ is already the central bridge from dynamics to rate.

One often writes

$$\overline{|\mathcal{M}|^2} = \frac{1}{N_{\text{in}}} \sum_{\text{spins, colours}} |\mathcal{M}|^2 , \tag{6.1}$$

where N_{in} is the number of initial-state spin and colour configurations over which one averages in an unpolarised calculation. The sum in the numerator runs over all final states and over any unobserved initial-state quantum numbers.

6.2 Initial-state averages and final-state sums

In a typical collider setting the incoming beams are not prepared in a fixed spin state, so one averages over the allowed initial polarisations. Final-state spins and colours that are not explicitly measured are summed over. This is why one distinguishes between $|\mathcal{M}|^2$ for a particular microscopic configuration and the experimentally more useful unpolarised quantity $\overline{|\mathcal{M}|^2}$.

For spin-$\frac{1}{2}$ fermions, the standard completeness relations are

$$\sum_s u(p,s)\bar{u}(p,s) = \gamma^\mu p_\mu + m, \qquad \sum_s v(p,s)\bar{v}(p,s) = \gamma^\mu p_\mu - m. \tag{6.2}$$

These identities are one reason why trace methods become useful in explicit calculations, although the present module uses them only selectively.

6.3 Polarisation and spin remarks at course level

When external vector bosons are involved, one must also sum over physical polarisation states. For scalar external particles no such sum is needed. For spin-$\frac{1}{2}$ particles, helicity or chirality structure can become phenomenologically important, especially in weak processes. This is one reason why angular distributions and current structure are so informative.

At the level of Module 6, the pedagogical point is not to develop a full spinor-technology course, but to make clear why spin information appears at all. It appears because the amplitude knows about the Lorentz structure of the interaction, and the Lorentz structure depends on the spin and current content of the external particles.

6.4 A bounded comment on what is deferred to advanced QFT courses

In advanced quantum field theory one learns systematic trace technology, helicity methods, gauge-invariant organisation of amplitudes, loop corrections, and more sophisticated treatments of external-state polarisation. These methods are powerful and essential in research. But the present module only needs the conceptual idea: the amplitude must be squared and appropriately summed or averaged before it becomes a rate.

Guided checks

- Can you explain why amplitudes are added before squaring?
- Can you explain why one averages over initial spins in an unpolarised beam?
- Can you explain why the final answer relevant for rates is often $\overline{|\mathcal{M}|^2}$ rather than $|\mathcal{M}|^2$ for one specific spin configuration?

7 Relativistic kinematics and Lorentz-invariant phase space

7.1 Four-momentum conservation in decays and scattering

Every decay and scattering process is constrained by four-momentum conservation. For a decay $A \to 1 + 2 + \cdots$, one has

$$p_A = p_1 + p_2 + \cdots . \tag{7.1}$$

For a scattering process $1 + 2 \to 3 + 4 + \cdots$, the same conservation law holds. This simple fact strongly shapes which channels are kinematically allowed and how the final-state momenta can be distributed.

In practice, a great deal of phenomenology is already determined by this constraint. Thresholds, invariant-mass peaks, endpoints, and allowed angular regions all reflect the underlying conservation laws.

7.2 Mandelstam variables at a first orientation level

For $2 \to 2$ scattering it is natural to introduce the Mandelstam variables

$$s = (p_1 + p_2)^2, \qquad t = (p_1 - p_3)^2, \qquad u = (p_1 - p_4)^2. \tag{7.2}$$

These variables are Lorentz invariant and therefore provide a convenient language for amplitudes and cross sections. The variable s measures the centre-of-mass energy squared available to the process, whereas t and u encode momentum transfer and angular information.

At introductory level, one should not treat the Mandelstam variables as a formal trick. They are precisely the invariants that allow us to describe scattering in a frame-independent way.

7.3 Phase space in concept

Even after the dynamics of a process have been specified by the amplitude, the final state may be realised in many different kinematic configurations. Phase space is the measure of these allowed configurations. It is therefore the kinematic complement to the dynamical information carried by $\mathcal{M}$.

The Lorentz-invariant phase-space element for a final state f is

$$d\Phi_f = (2\pi)^4 \delta^{(4)}\left(p_{\text{in}} - \sum_i p_i\right) \prod_i \frac{d^3 p_i}{(2\pi)^3\, 2E_i}. \tag{7.3}$$

A useful way to think about phase space is that it answers the question: after the interaction has occurred, how many kinematically allowed ways are there to distribute the final-state momenta while respecting energy-momentum conservation? The rate depends both on how strongly the process occurs and on how much kinematic room the final state has.

7.4 Why phase-space factors matter physically

Phase space is not a decorative measure factor. It explains many familiar physical patterns. Decays close to threshold are suppressed because the final state has little momentum room. Multibody final states can be dynamically possible but still relatively suppressed because the phase-space integration is more involved. Resonance line shapes, threshold behaviour, and the contrast between two-body and many-body channels all become easier to understand once the role of phase space is kept in view.

7.5 Two-body final states as the cleanest starting point

Two-body decays and two-body scattering processes are the cleanest pedagogical starting point because much of the kinematics can be made explicit. In a two-body decay in the rest frame of the parent particle, the final-state particles emerge back-to-back with fixed magnitudes of momentum. This makes visible the relation between the masses, the available energy, and the resulting width.

For a decay $A \to 1 + 2$ in the rest frame of A, the common magnitude of the daughter three-momenta is

$$|\mathbf{p}| = \frac{\lambda^{1/2}(m_A^2, m_1^2, m_2^2)}{2m_A}, \tag{7.4}$$

where the Källén function is

$$\lambda(x, y, z) = x^2 + y^2 + z^2 - 2xy - 2xz - 2yz. \tag{7.5}$$

The corresponding two-body phase-space element becomes

$$d\Phi_2 = \frac{1}{16\pi^2} \frac{|\mathbf{p}|}{m_A} \, d\Omega \qquad (1 \to 2 \text{ in the parent rest frame}). \tag{7.6}$$

For a $2 \to 2$ process in the centre-of-mass frame, one similarly finds

$$d\Phi_2 = \frac{1}{16\pi^2} \frac{|\mathbf{p}_f|}{\sqrt{s}} \, d\Omega. \tag{7.7}$$

This is why so many clean textbook formulas reduce to a simple angular distribution multiplied by a momentum ratio.

Remark 7.1: Dynamics and kinematics are different ingredients

The amplitude tells us how strongly the process occurs; the phase-space measure tells us how the kinematics of the final state are distributed. A good physical understanding of a process requires both ingredients.

8 Particle decays, decay widths, and lifetimes

8.1 What it means for an unstable particle to decay

An unstable particle is one for which the theory allows a transition to lighter final states consistent with the relevant conservation laws. The weak interaction plays a particularly important role here because it changes flavour and therefore opens channels that would not be accessible through strong or electromagnetic interactions alone. Many observed particles are therefore not stable objects but resonant or metastable excitations that decay after some characteristic lifetime.

From the phenomenological point of view, a decay is the simplest predictive process because it begins with a single parent state. Yet even here the same core logic applies: the interaction structure determines the amplitude, the squared amplitude must be combined with phase space, and the result is the decay width.

8.2 Total decay width and lifetime

The differential decay rate of a parent particle A is

$$d\Gamma(A \to f) = \frac{1}{2m_A}\,\overline{|\mathcal{M}|^2}\,d\Phi_f. \tag{8.1}$$

This relation is the relativistic field-theory analogue of Fermi's golden rule: dynamics and density of final states appear together. The total decay width Γ_A of a particle A is the sum of the partial widths into all accessible final states. It sets the overall decay probability per unit time and is related to the lifetime by

$$\tau_A = \frac{1}{\Gamma_A}. \tag{8.2}$$

At course level, one may think of the width as the inverse timescale of the instability. Narrow states live relatively long and tend to produce sharp resonance features; broad states are short-lived and correspond to less sharply defined invariant-mass structures.

8.3 Partial decay widths

If the particle can decay into several distinct final states, one defines a partial width Γ_i for each channel. The total width is then

$$\Gamma_{\text{tot}} = \sum_i \Gamma_i. \tag{8.3}$$

This decomposition is central experimentally because different channels have different cleanliness, different backgrounds, and different theoretical interest. The decay pattern of a particle therefore carries information not just about its existence but about its couplings and quantum numbers.

8.4 Two-body decay kinematics

For a two-body decay $A \to 1 + 2$, the general formula (8.1) becomes

$$d\Gamma_{1\to 2} = \frac{1}{32\pi^2}\,\frac{|\mathbf{p}|}{m_A^2}\,\overline{|\mathcal{M}|^2}\,d\Omega, \tag{8.4}$$

with

$$|\mathbf{p}| = \frac{\lambda^{1/2}(m_A^2, m_1^2, m_2^2)}{2m_A}. \tag{8.5}$$

If the squared amplitude is independent of the decay angles, or once the angular integration has been carried out, one obtains the standard result

$$\Gamma(A \to 1 + 2) = \frac{|\mathbf{p}|}{8\pi m_A^2}\,\overline{|\mathcal{M}|^2} = \frac{\lambda^{1/2}(m_A^2, m_1^2, m_2^2)}{16\pi m_A^3}\,\overline{|\mathcal{M}|^2}. \tag{8.6}$$

This formula is one of the workhorses of introductory phenomenology. It makes visible that a decay width can be small either because the coupling is small or because the available phase space is small.

Example 8.1: A generic two-body decay in constant-amplitude approximation

Suppose that for a decay $A \to 1 + 2$ the angle-integrated squared amplitude is approximately constant. Then (8.6) shows immediately that the width vanishes at threshold through the factor $\lambda^{1/2}$. This is the cleanest first demonstration that even a nonzero coupling does not guarantee a large width: kinematics can suppress the channel strongly.

8.5 Why decay widths are observable bridges between theory and experiment

The decay width is one of the most direct bridges between formal theory and data. It is calculable from the interaction structure, yet it has immediate empirical meaning through lifetimes, branching fractions, resonance widths, and visible decay modes. A decay calculation therefore exemplifies the whole phenomenological programme of the module in concentrated form.

A useful Standard-Model illustration is the leptonic decay of the W boson. Neglecting final-state masses, the tree-level result is

$$\Gamma(W \to \ell\nu_\ell) = \frac{G_F m_W^3}{6\sqrt{2}\,\pi}, \tag{8.7}$$

which shows in a compact way how the weak scale and the weak coupling set a measurable partial width. Even when one does not derive this formula line by line, its structure is pedagogically informative: the width is controlled by the interaction strength and by the fact that a heavy vector boson has a large phase space for decay.

Example 8.2: Weak decays as structurally informative processes

Muon decay, beta decay, and gauge-boson decays are especially instructive because they reveal the chiral and flavour-changing structure of the weak interaction. Even when treated only at introductory level, they show how a specific interaction structure leaves a measurable imprint on lifetimes and final states.

9 Branching ratios and decay chains

9.1 Definition of branching ratio

The branching ratio of a channel i is defined by

$$\mathrm{BR}_i = \frac{\Gamma_i}{\Gamma_{\mathrm{tot}}}. \tag{9.1}$$

It answers a very practical question: when the parent particle decays, how often does it do so through this specific final state? Since experiments identify particles through channels, not through abstract total widths alone, branching ratios are often the more operationally useful quantity. By definition one has

$$\sum_i \mathrm{BR}_i = 1. \tag{9.2}$$

9.2 Why branching ratios are often more useful than total widths alone

A total width tells us how unstable a particle is overall, but it does not tell us which final states dominate or which channel is experimentally most accessible. Branching ratios encode this composition. They are therefore indispensable in both theory and experiment. A channel with a small branching ratio may still be the preferred discovery mode if its background is exceptionally low, whereas a channel with a large branching ratio may be difficult to exploit if the background is overwhelming.

9.3 Dominant versus rare channels

The hierarchy among branching ratios reflects both dynamics and kinematics. Strong decays, when allowed, are often much faster than electromagnetic ones, and electromagnetic decays are often faster than weak ones. But there are also kinematic and symmetry effects that can suppress a channel even when it is not forbidden. For this reason decay patterns are rich diagnostic tools.

At introductory level the key lesson is not to memorise lists of channels, but to understand that branching ratios reveal coupling structure, phase-space availability, and sometimes selection rules.

9.4 Decay chains and experimentally visible final states

A detector often does not observe the parent particle directly. Instead it observes a decay chain, possibly with intermediate states that themselves decay. This makes branching ratios multiplicative along a decay chain:

$$\mathrm{BR}(A \to B\,C \to f_B\,f_C) = \mathrm{BR}(A \to B\,C)\,\mathrm{BR}(B \to f_B)\,\mathrm{BR}(C \to f_C), \tag{9.3}$$

provided the intermediate states are treated in the usual narrow-width picture. This explains why experimentally visible final states may differ significantly from the most elementary theoretical channel language.

For this reason phenomenology must always keep one eye on the immediate process and another on the visible end products. The event yield in a specific channel is often of the schematic form

$$N_f \sim \mathcal{L}_{\mathrm{int}}\,\sigma(\text{production})\,\mathrm{BR}(\text{decay chain})\,A\,\epsilon, \tag{9.4}$$

where A and ϵ represent acceptance and efficiency. This is already the beginning of the bridge toward analysis-level thinking.

10 Scattering cross sections

10.1 Physical meaning of a cross section

A cross section is the standard measure of the effective interaction rate in scattering. It is not a literal geometric area in any naive classical sense, although historically the term suggests one. Rather, it is the quantity that relates the incoming flux of particles to the rate of occurrence of a given final state.

Because collider experiments are built around controlled initial states and measured event yields, the cross section is one of the central measurable outputs of Standard Model phenomenology. In a practical sense, if one knows the luminosity and the relevant efficiencies, the cross section is what turns theory into an expected event count.

10.2 Flux factor and interaction probability

The general Lorentz-invariant structure of a scattering cross section is

$$d\sigma = \frac{1}{F}\,\overline{|\mathcal{M}|^2}\,d\Phi_f, \tag{10.1}$$

where F is the appropriate incoming flux factor. In invariant form one may write

$$F = 4\sqrt{(p_1 \cdot p_2)^2 - m_1^2 m_2^2}. \tag{10.2}$$

The flux factor is the scattering counterpart of the $1/(2m_A)$ factor in decays. It accounts for the kinematic normalisation of the incoming state.

For $2 \to 2$ scattering in the centre-of-mass frame one obtains the standard result

$$\frac{d\sigma}{d\Omega} = \frac{1}{64\pi^2 s}\,\frac{|\mathbf{p}_f|}{|\mathbf{p}_i|}\,\overline{|\mathcal{M}|^2}. \tag{10.3}$$

When both initial and final masses may be neglected, the momentum ratio is unity and the formula simplifies even further.

10.3 Total versus differential cross sections

The total cross section integrates over the full allowed final-state kinematics. A differential cross section keeps one or more kinematic variables unintegrated and therefore reveals more detail. For example, $d\sigma/d\Omega$ gives angular information, while $d\sigma/dm$ can expose a resonance structure.

Differential information is often where the real physics becomes visible. Two models may predict similar total rates but very different distributions. This is why modern collider phenomenology places enormous weight on differential observables.

10.4 Why cross sections are central in collider physics

In collider physics the cross section is the natural interface between microscopic dynamics and macroscopic counting. It is what allows one to predict how often a channel should appear for a given integrated luminosity. Even a schematic relation such as

$$N_{\text{events}} \sim \mathcal{L}_{\text{int}}\,\sigma\ \text{(branching ratios) (acceptance) (efficiency)} \tag{10.4}$$

already makes clear why cross sections occupy such a central place in experimental interpretation. A cross section is therefore not only a formal quantity; it is the number that connects the hard-scattering theory to what an experiment may realistically record.

Take-home message

A decay width characterises how a state disappears. A cross section characterises how frequently a specified final state is produced from a given initial state. Both are derived from the same core ingredients: amplitudes, squared amplitudes, and kinematics.

11 Differential cross sections and selected observables

11.1 Angular distributions

Angular distributions are among the simplest differential observables and often reveal the spin and current structure of the interaction. At tree level, the dependence of $d\sigma/d\Omega$ on the scattering angle can already distinguish different mediators or different Lorentz structures. This is one reason why differential distributions often carry more structural information than total rates alone.

A classic example is the QED process $e^+e^- \to \mu^+\mu^-$. In the relativistic limit the distribution behaves as $1 + \cos^2\theta$, which is far from isotropic and already encodes the vector-current nature of the interaction. Such examples make clear why the shape of a distribution is often as important as its normalisation.

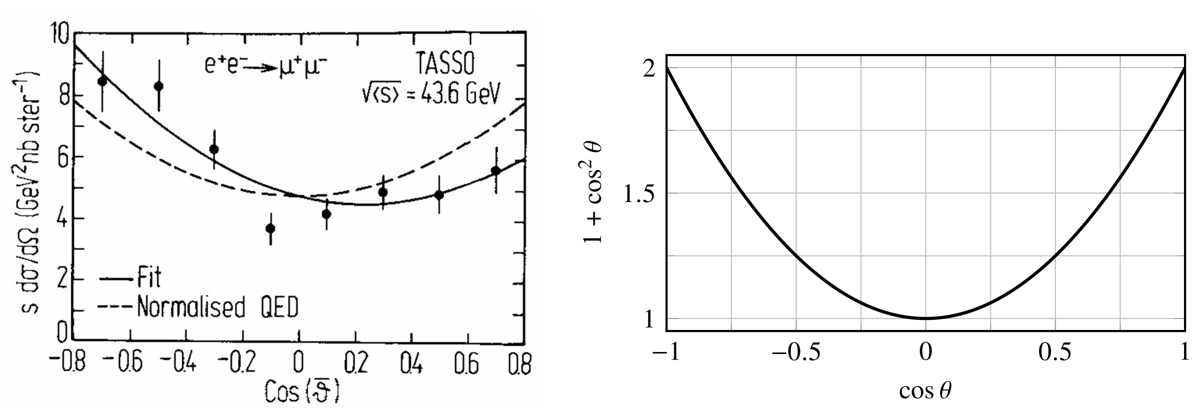

Figure 3: Left: experimental measurement of the angular distribution for $e^+e^- \to \mu^+\mu^-$ from the TASSO experiment at $\sqrt{s} = 43.6$ GeV, showing clear evidence that the distribution is not isotropic and is broadly consistent with the characteristic vector-current behaviour discussed in the text. Right: internal pedagogical plot of the idealised relativistic-QED shape proportional to $1 + \cos^2\theta$. The purpose of placing the two panels side by side is to connect the schematic theoretical expectation to real experimental evidence.

11.2 Invariant mass, transverse momentum, and rapidity / pseudorapidity as orientation-level observables

The invariant mass of a subsystem,

$$m_{ab}^2 = (p_a + p_b)^2, \tag{11.1}$$

is one of the most important collider observables because it can reveal intermediate resonant states. The transverse momentum p_T is especially useful at hadron colliders because the beam direction singles out a preferred axis and the transverse plane is where conservation laws are experimentally most robust.

For collider orientation it is also useful to define rapidity and pseudorapidity,

$$y = \frac{1}{2} \ln \frac{E + p_z}{E - p_z}, \qquad \eta = -\ln \tan \frac{\theta}{2}. \tag{11.2}$$

For highly relativistic particles y and η are close numerically. At the level of Module 6, these quantities are introduced as the natural language in which theoretical processes are compared to measured distributions. Detailed detector-level treatment is deferred.

11.3 Resonances and peaks in invariant-mass distributions

A resonance is often first seen experimentally as an enhancement or peak in an invariant-mass distribution. This is a particularly striking illustration of the bridge between amplitudes and observables. The propagator structure of an intermediate state leaves a direct imprint on the measured distribution. Width, mass, and sometimes interference all become visible in the line shape.

This is one reason why invariant-mass plots are so central in particle physics. They are often the place where the theory announces the presence of a particle most directly.

11.4 Why distributions reveal dynamics beyond a single total rate

A total rate compresses a great deal of information into one number. Distributions recover some of that lost information. They can reveal thresholds, resonance structures, spin effects, and phase-space boundaries. In modern phenomenology, the move from integrated to differential observables is therefore not a luxury; it is often what makes the physics discernible.

12 Canonical tree-level examples

12.1 Electron–positron annihilation as a prototype scattering process

The process

$$e^+(p_2)\, e^-(p_1) \to \mu^-(k_1)\, \mu^+(k_2) \tag{12.1}$$

is one of the cleanest textbook examples because, at lowest order in QED and away from the Z pole, it is described by a single s-channel photon-exchange diagram. The basic topology is shown in Fig. 4. Using the QED vertex and photon propagator, one finds

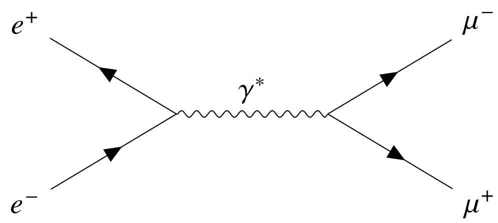

Figure 4: Tree-level QED diagram for $e^+e^- \to \mu^+\mu^-$. Away from the Z pole this single s-channel photon-exchange diagram is the cleanest first example of how a Standard Model interaction vertex, combined with a propagator, becomes a calculable amplitude and then a cross section.

$$\mathrm{i}\mathcal{M} = [\bar{v}(p_2)(-\mathrm{i}e\gamma^\mu)u(p_1)]\,\frac{-\mathrm{i}g_{\mu\nu}}{s}\,[\bar{u}(k_1)(-\mathrm{i}e\gamma^\nu)v(k_2)]\,. \tag{12.2}$$

Equivalently,

$$\mathcal{M} = \frac{e^2}{s}\,[\bar{v}(p_2)\gamma^\mu u(p_1)]\,\left[\bar{u}(k_1)\gamma_\mu v(k_2)\right], \tag{12.3}$$

up to overall sign conventions. This already exhibits the universal grammar of tree-level phenomenology: a vertex, a propagator, and another vertex.

In the relativistic limit $m_e, m_\mu \ll \sqrt{s}$, a standard spin-sum calculation gives

$$\overline{|\mathcal{M}|^2} = e^4\left(1 + \cos^2\theta\right), \tag{12.4}$$

where θ is the centre-of-mass scattering angle of the outgoing muon. Inserting this into (10.3) yields

$$\frac{\mathrm{d}\sigma}{\mathrm{d}\Omega} = \frac{\alpha^2}{4s}\left(1 + \cos^2\theta\right), \tag{12.5}$$

and after integration over solid angle,

$$\sigma(e^+e^- \to \mu^+\mu^-) = \frac{4\pi\alpha^2}{3s}. \tag{12.6}$$

This example is pedagogically valuable because the relation between diagram, amplitude, angular distribution, and total cross section is unusually transparent. It also teaches that the cross section falls as $1/s$ at high energy in pure QED. In the full Standard Model, once $\sqrt{s}$ approaches m_Z, Z exchange and γ–Z interference must be included, and the LEP Z resonance becomes visible in the same channel.

12.2 Muon decay as a prototype weak decay process

Muon decay is the corresponding prototype on the decay side:

$$\mu^- \to e^-\,\bar{\nu}_e\,\nu_\mu. \tag{12.7}$$

At the Standard Model level it proceeds through W exchange and therefore provides a direct first look at weak charged-current phenomenology. Because the muon is a lepton, the process is conceptually clean: strong-interaction complications are absent.

At momentum transfers much smaller than m_W^2, the propagator of the exchanged W boson may be replaced by a local effective interaction. The full Standard Model topology and the low-energy effective picture are shown side by side in Fig. 5. One writes

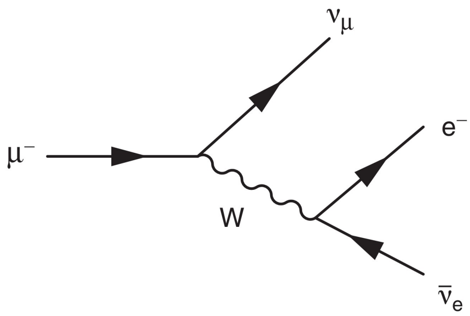

Tree-level Standard Model W-exchange diagram for $\mu^- \to e^- \bar{\nu}_e \nu_\mu$.

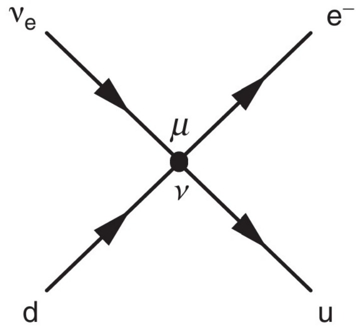

Illustration of the low-energy local charged-current four-fermion interaction used to visualise the effective Fermi-theory limit.

Figure 5: Muon decay as an instructive comparison between the full Standard Model weak-interaction description and its low-energy effective interpretation. The left panel shows the mediator-based W-exchange picture for $\mu^- \to e^- \bar{\nu}_e \nu_\mu$. The right panel shows a textbook representation of the local charged-current four-fermion interaction that emerges in the regime $q^2 \ll m_W^2$. In the present subsection, the second panel is used only to visualise the effective contact-interaction limit of weak decays.

$$\mathcal{L}_F = -\frac{G_F}{\sqrt{2}} \left[\bar{e}\,\gamma^\mu(1-\gamma^5)\nu_e\right] \left[\bar{\nu}_\mu\,\gamma_\mu(1-\gamma^5)\mu\right]. \tag{12.8}$$

This is already enough to see that the decay is controlled by a four-fermion interaction, by the chiral $(V - A)$ structure of the weak current, and by a three-body phase-space integration. The latter is precisely why the muon decay calculation is richer than a two-body example.

A standard calculation, neglecting the electron mass and radiative corrections, gives

$$\Gamma_\mu \simeq \frac{G_F^2 m_\mu^5}{192\pi^3}. \tag{12.9}$$

One may also write the squared amplitude in a useful schematic form as

$$\overline{|\mathcal{M}|^2} = 64\, G_F^2\, (p_\mu \cdot p_{\bar{\nu}_e})\, (p_e \cdot p_{\nu_\mu}), \tag{12.10}$$

again in the limit of negligible electron mass. The full width then follows from integrating this over three-body phase space. This result is physically important because it shows how a weak coupling and a large power of the muon mass combine to produce a measurable lifetime. In fact, precision measurement of the muon lifetime is one of the classic routes to determining G_F.

12.3 A bounded low-energy weak-interaction remark: the 4-Fermi limit

The 4-Fermi limit is obtained by expanding the W propagator for $q^2 \ll m_W^2$:

$$\frac{-g_{\mu\nu}}{q^2 - m_W^2} \approx \frac{g_{\mu\nu}}{m_W^2}. \tag{12.11}$$

This converts the mediator-based Standard Model process into an effective local interaction. Matching the resulting coefficient to (12.8) gives the standard relation

$$\frac{G_F}{\sqrt{2}} = \frac{g^2}{8m_W^2}. \tag{12.12}$$

This is a beautiful example of effective-field-theory logic inside Standard Model phenomenology: the heavy mediator is still there in the full theory, but at low energy the process is described economically by a local operator.

12.4 What these examples teach conceptually

The purpose of these examples is not to burden the module with lengthy algebra. It is to make visible the common pattern. In e^+e^- annihilation and in muon decay alike, one begins with an interaction structure, constructs the relevant diagram, writes the amplitude, squares it appropriately, and then combines it with kinematics and phase space. The same logic then generalises to more complicated Standard Model processes.

Example 12.1: Why these two examples are especially useful

Electron–positron annihilation is a clean prototype for scattering because the initial state is simple and the lowest-order diagram is transparent. Muon decay is a clean prototype for weak decays because it isolates charged-current weak phenomenology without strong-interaction complications. Together they illustrate both sides of the module: production and decay.

13 Resonances, propagators, and unstable intermediate states

13.1 Internal propagators and kinematic enhancement

Internal propagators contribute denominators of the schematic form

$$\frac{1}{q^2 - m_X^2} \tag{13.1}$$

for an exchanged particle of mass m_X. When the kinematics of the process bring q^2 near m_X^2, the contribution can be enhanced and the process may exhibit resonant behaviour. This is one of the most direct ways in which the analytic structure of the amplitude becomes visible in a measurable observable.

13.2 Resonance idea at course level

At course level, a resonance may be understood as the experimental manifestation of an unstable intermediate state whose propagator becomes important in a particular invariant-mass region. Because the state is unstable, one replaces the simple pole by the width-regulated expression

$$\frac{1}{s - m_X^2 + \mathrm{i}m_X\Gamma_X}. \tag{13.2}$$

After squaring the amplitude, this leads to a Breit–Wigner-type structure

$$|\mathcal{M}|^2 \propto \frac{1}{(s - m_X^2)^2 + m_X^2 \Gamma_X^2}, \tag{13.3}$$

up to numerator factors and possible interference with non-resonant contributions.

13.3 Narrow-width interpretation at qualitative level

If $\Gamma_X \ll m_X$, the resonance is narrow and its mass can often be read off rather cleanly from an invariant-mass peak. This is the qualitative content of the narrow-width picture. It helps explain why some particles appear as sharp resonances while others are much broader and less cleanly resolved.

A mathematically useful form of the same idea is the narrow-width approximation,

$$\frac{1}{(s - m_X^2)^2 + m_X^2 \Gamma_X^2} \longrightarrow \frac{\pi}{m_X \Gamma_X} \delta(s - m_X^2) \qquad (\Gamma_X / m_X \ll 1). \tag{13.4}$$

In this regime production and decay often factorise approximately:

$$\sigma(ab \to X \to f) \approx \sigma(ab \to X)\, \mathrm{BR}(X \to f). \tag{13.5}$$

This approximation is not exact in all cases, but it is extremely useful and physically transparent.

13.4 Why resonances are experimentally powerful

Resonances are powerful experimentally because they concentrate the signal into a characteristic kinematic region. This enhances interpretability and often allows a particle to be identified through its decay products even if the particle itself is too short-lived to be observed directly. Much of collider spectroscopy, old and new, relies on precisely this logic.

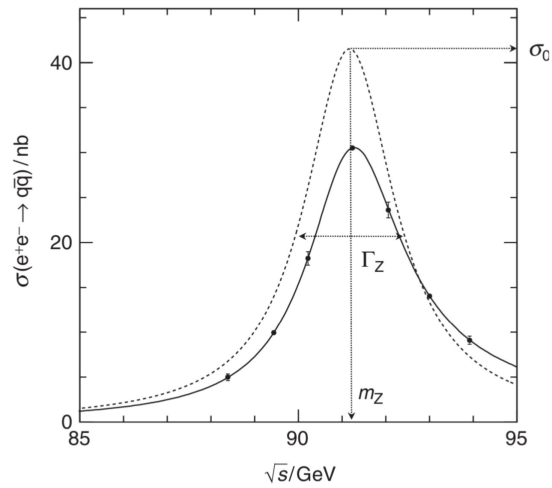

Figure 6: Experimental illustration of the Z-resonance line shape in $e^+e^- \to q\bar{q}$ as a function of the centre-of-mass energy $\sqrt{s}$. The plot makes visible the central phenomenological meanings of the resonance mass m_Z, the total width Γ_Z, and the peak cross section σ_0. Such figures provide direct experimental evidence that unstable intermediate states appear as characteristic enhancements in invariant-energy distributions, with the resonance position identifying the particle mass and the line shape encoding its finite lifetime through the width.

Example 13.1: A classic Standard Model resonance picture

The LEP Z-pole programme provides a textbook example. In $e^+e^- \to f\bar{f}$, the cross section is dramatically enhanced when $\sqrt{s}$ is tuned close to m_Z. The position of the peak measures the mass, the line shape encodes the width, and the decay channels reveal the couplings and branching ratios.

14 From parton-level processes to collider-facing observables

14.1 Why hadron-collider interpretation is richer than textbook parton scattering

Textbook processes such as $e^+e^- \to \mu^+\mu^-$ are deliberately clean. Hadron-collider processes are richer because the incoming protons are composite. One therefore does not collide fixed elementary initial states in the same sense, but rather partons drawn from parton distributions inside the proton. The hard process is then embedded inside a broader hadronic environment.

This additional layer does not invalidate the basic phenomenological logic of the module. It enriches it. One still computes parton-level amplitudes and cross sections, but these are folded with information about the proton structure and then related to visible final states.

14.2 Partons, visible final states, and reconstruction-level observables

At orientation level, the relation between hadron-level and parton-level physics may be summarised schematically as

$$\sigma(pp \to X) = \sum_{ij} \int \mathrm{d}x_1\, \mathrm{d}x_2\, f_i(x_1, \mu_F) f_j(x_2, \mu_F)\, \hat{\sigma}_{ij \to X}(x_1 x_2 s, \mu_F, \mu_R). \tag{14.1}$$

The partonic cross section $\hat{\sigma}$ is the direct output of the hard-scattering calculation; the measured final state, however, is built from reconstructed jets, leptons, photons, and missing transverse momentum. This is why collider phenomenology always involves a careful translation between idealised parton-level reasoning and experimentally accessible quantities.

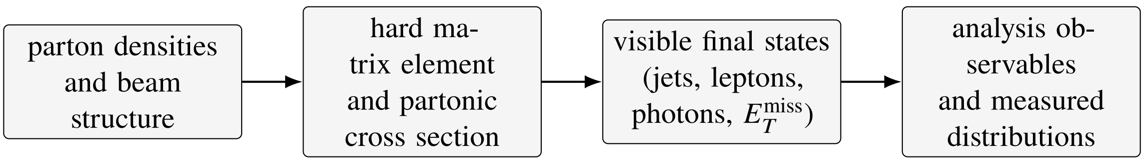

Figure 7: A deliberately simplified workflow showing how hadron-collider predictions are layered. The point is not to replace a full collider-phenomenology treatment, but to help the reader see why partonic calculations, visible final states, and analysis observables are related but not identical.

14.3 Why theory predictions and measured distributions must be connected carefully

Even when the hard process is understood, the experiment measures selected events after detector response, reconstruction, and analysis choices. For that reason phenomenology must respect a layered structure: hard matrix element, event kinematics, visible final states, and finally analysis observables. Module 6 only begins this story, but it is important that the student already sees why the connection is careful rather than automatic.

14.4 Pedagogical boundary: detailed detector and analysis issues belong later

A full treatment of detector response, object identification, backgrounds, unfolding, and statistical inference belongs to later laboratory or analysis-focused work. The present section is intentionally a bridge. Its task is to make clear that the route from theory to data is real and structured, without turning Module 6 into a full experimental manual.

15 What Module 6 teaches us about testing the Standard Model

15.1 From formal theory to falsifiable predictions

The Standard Model becomes scientifically powerful not merely because it is mathematically elegant, but because it yields falsifiable predictions. Decay widths, cross sections, line shapes, and differential distributions are among the clearest forms of those predictions. They can be measured, compared to the theory, and used either to confirm the Standard Model or to reveal its limitations.

15.2 Why decay rates and cross sections are structural observables

Decay rates and cross sections are structural observables because they sit directly at the interface of dynamics and experiment. They depend on couplings, masses, current structure, phase space, and kinematic selection. They are therefore especially sensitive to the inner logic of the theory. A disagreement in a measured width or cross section is not a peripheral issue; it is often a direct challenge to the assumed interaction structure.

15.3 Why precision and distributions matter

As experiments become more precise, total rates alone are no longer enough. Distributions, correlations, and differential measurements become increasingly central. They test not only whether a channel exists, but whether the detailed dynamics are correctly described. This is one reason why modern phenomenology continuously moves toward more differential and more precise observables.

15.4 The Standard Model as a predictive framework, not a particle catalogue

A recurrent theme of the course is that the Standard Model should be understood as a constrained theoretical framework rather than as a list of particles to memorise. Module 6 sharpens this point. The theory predicts not just which particles exist, but how often they are produced, how they decay, and how their kinematic distributions should look. This predictive power is what makes the Standard Model scientifically compelling.

Take-home message

The Standard Model is tested through observables. Widths, branching ratios, cross sections, and differential distributions are not secondary outputs; they are the central predictive quantities through which the formal theory confronts experiment.

16 A bounded outlook: what Module 6 will not attempt

16.1 Full loop calculations and renormalisation

Higher-order corrections and renormalisation are essential in precision phenomenology, but they require a machinery beyond the scope of the present note. Module 6 signposts their importance without making them core material. The priority here is the tree-level logic that makes the predictive structure visible.

16.2 Advanced collider simulation and parton showers

Modern collider predictions involve parton showers, hadronisation models, matching, merging, and event simulation frameworks. These are indispensable in research-level collider phenomenology, especially at hadron colliders. Yet they presuppose that the student already understands the matrix-element level and the role of observables. That is why they lie beyond the present module.

16.3 Detector systematics and full statistical inference

Experimental interpretation also requires background modelling, detector systematics, efficiency corrections, and statistical inference. These belong naturally to analysis-oriented or laboratory-focused study. Module 6 prepares the conceptual ground by clarifying what quantity the theory predicts and why that quantity matters, but it does not attempt the whole data-analysis chain.

16.4 Why these belong to later or more advanced study

The boundaries of the module are pedagogical rather than dismissive. Nothing listed above is unimportant. On the contrary, these topics become more meaningful once the student can already explain the role of amplitudes, widths, cross sections, and distributions. The present note therefore treats them as the next layer rather than as the starting point.

17 Bridge to later modules

17.1 From phenomenology to limitations and open questions: Module 7

Once the Standard Model has been expressed in terms of measurable observables, its limits become easier to formulate sharply. Are the predicted rates correct in every channel? Do flavour observables reveal tensions? Are there sectors, such as neutrino masses or dark matter, where the minimal theory has no satisfactory account? Module 7 takes up precisely these questions. Module 6 therefore provides the predictive language in which many later shortcomings are recognised.

17.2 From observables to analysis workflows: Module 8

Module 8 moves from formal phenomenology toward practical data analysis. Histograms, distributions, event selection, simple fitting, and comparison of predicted and observed yields all presuppose the language introduced here. The observables of Module 6 become the analysis objects of Module 8. In that sense Module 6 is not only a bridge from theory to experiment; it is also the conceptual precondition for the practical laboratory and project component.

17.3 From theoretical prediction to data interpretation

Seen in the full course structure, the sequence is now clear. Modules 1–5 construct the theory. Module 6 translates it into phenomenological predictions. Module 7 asks where the predictive framework is incomplete. Module 8 turns the language of observables into concrete analysis practice. This sequence is one of the reasons why the present module occupies such a structurally important position.

18 Final summary and conceptual map

18.1 The logical chain of Module 6

The core logic of the module can be summarised succinctly. The Standard Model Lagrangian determines the allowed interaction terms. These interaction terms determine vertices and propagators. Feynman rules translate these into amplitudes for processes. Different contributions must be added at amplitude level, after which the squared amplitude is combined with spin sums or averages, phase space, and the appropriate decay or flux normalisation. The resulting quantities are decay widths, branching ratios, cross sections, and differential observables. These are the quantities that experiments measure and compare with theory.

18.2 What the student should now be able to see

After this module, the student should be able to explain what Feynman diagrams encode, why amplitudes rather than probabilities are added, why $|\mathcal{M}|^2$ controls rates, why phase space matters, how decay widths differ from branching ratios, how cross sections differ from event counts, and why invariant masses and differential distributions are so central in particle-physics experiments. The student should also be able to recognise why simple tree-level examples are not childish toy models but the cleanest first realisations of the theory-to-observable bridge.

18.3 From amplitudes to observables to data

The broader message of Module 6 is that the Standard Model does not end with symmetries and couplings. Its real predictive force appears when those structures are translated into measurable quantities. Phenomenology is therefore not the abandonment of theory but its operational completion. Through decays, cross sections, and observables, the Standard Model becomes a framework that can be tested, refined, and, where necessary, challenged by data.

Take-home message

Module 6 turns the Standard Model into a predictive experimental language. Interaction terms lead to Feynman rules, rules lead to amplitudes, amplitudes lead to widths and cross sections, and these lead to the observables through which the theory confronts data. In that sense phenomenology is the bridge between formal structure and physical evidence.

A Conventions and notation summary

Quantity	Convention
Metric	$\eta_{\mu\nu} = \mathrm{diag}(1, -1, -1, -1)$
Natural units	$\hbar = c = 1$
Amplitude	$\mathcal{M}$ or $\mathcal{A}$
Total decay width	Γ
Lifetime	$\tau = 1/\Gamma$
Branching ratio	$\mathrm{BR}_i = \Gamma_i/\Gamma_{\mathrm{tot}}$
Cross section	σ, with differential forms such as $\mathrm{d}\sigma/\mathrm{d}\Omega$
Mandelstam variables	$s = (p_1 + p_2)^2$, $t = (p_1 - p_3)^2$, $u = (p_1 - p_4)^2$
Selected collider variables	p_T, η, m_{ab}

B Feynman-rule and amplitude workflow summary

The conceptual workflow of a tree-level process may be summarised as follows:

$$
\begin{aligned}
\mathcal{L}_{\mathrm{SM}} &\longrightarrow \mathcal{L}_{\mathrm{free}} + \mathcal{L}_{\mathrm{int}} && \text{(B.1)}\\
&\longrightarrow \text{free propagators and interaction vertices} && \text{(B.2)}\\
&\longrightarrow \text{Feynman diagrams for the chosen process} && \text{(B.3)}\\
&\longrightarrow \mathcal{M}_{\mathrm{tot}} = \sum_k \mathcal{M}_k && \text{(B.4)}\\
&\longrightarrow \overline{|\mathcal{M}_{\mathrm{tot}}|^2}. && \text{(B.5)}
\end{aligned}
$$

A few standard elements that repeatedly appear are summarised in Table 1.

Ingredient	Typical form
QED vertex	$-\mathrm{i}eQ_f\gamma^\mu$
Fermion propagator	$\dfrac{\mathrm{i}(\gamma^\mu q_\mu + m)}{q^2 - m^2 + \mathrm{i}\epsilon}$
Photon propagator (Feynman gauge)	$\dfrac{-\mathrm{i}g_{\mu\nu}}{q^2 + \mathrm{i}\epsilon}$
Amplitude sum	$\mathcal{M}_{\mathrm{tot}} = \sum_k \mathcal{M}_k$
Interference	$\lvert\mathcal{M}_1 + \mathcal{M}_2\rvert^2 = \lvert\mathcal{M}_1\rvert^2 + \lvert\mathcal{M}_2\rvert^2 + 2\Re(\mathcal{M}_1\mathcal{M}_2^*)$

Table 1: Representative building blocks in the tree-level workflow from Lagrangian to amplitude.

This is the minimal route from formal interaction structure to the quantity that enters a rate.

C Decay-width and branching-ratio summary

The generic Lorentz-invariant expression for a decay of a parent particle A is

$$\Gamma(A \to f) = \frac{1}{2m_A} \int d\Phi_f \, \overline{|\mathcal{M}|^2}. \tag{C.1}$$

For a two-body decay $A \to 1 + 2$ in the rest frame of A,

$$\Gamma(A \to 1 + 2) = \frac{|\mathbf{p}|}{8\pi m_A^2} \overline{|\mathcal{M}|^2} = \frac{\lambda^{1/2}(m_A^2, m_1^2, m_2^2)}{16\pi m_A^3} \overline{|\mathcal{M}|^2}, \tag{C.2}$$

where

$$|\mathbf{p}| = \frac{\lambda^{1/2}(m_A^2, m_1^2, m_2^2)}{2m_A}. \tag{C.3}$$

For several allowed channels,

$$\Gamma_{\text{tot}} = \sum_i \Gamma_i, \qquad \text{BR}_i = \frac{\Gamma_i}{\Gamma_{\text{tot}}}, \qquad \tau = \frac{1}{\Gamma_{\text{tot}}}. \tag{C.4}$$

A useful decay-chain relation is

$$\text{BR}(A \to B\,C \to f_B\, f_C) = \text{BR}(A \to B\,C)\,\text{BR}(B \to f_B)\,\text{BR}(C \to f_C). \tag{C.5}$$

A classic Standard Model three-body example is muon decay:

$$\Gamma_\mu \simeq \frac{G_F^2 m_\mu^5}{192\pi^3} \qquad (m_e \to 0, \text{ tree level}). \tag{C.6}$$

D Cross-section and phase-space summary

A generic scattering cross section has the form

$$d\sigma = \frac{1}{F} \overline{|\mathcal{M}|^2} \, d\Phi_f, \tag{D.1}$$

with Lorentz-invariant flux factor

$$F = 4\sqrt{(p_1 \cdot p_2)^2 - m_1^2 m_2^2}. \tag{D.2}$$

The Lorentz-invariant phase-space measure is

$$d\Phi_f = (2\pi)^4 \delta^{(4)}\left(p_{\text{in}} - \sum_i p_i\right) \prod_i \frac{d^3 p_i}{(2\pi)^3 2E_i}. \tag{D.3}$$

For $2 \to 2$ scattering in the centre-of-mass frame,

$$\frac{d\sigma}{d\Omega} = \frac{1}{64\pi^2 s} \frac{|\mathbf{p}_f|}{|\mathbf{p}_i|} \overline{|\mathcal{M}|^2}. \tag{D.4}$$

For a massless final state in a massless initial state this simplifies to

$$\frac{d\sigma}{d\Omega} = \frac{1}{64\pi^2 s}\overline{|\mathcal{M}|^2}. \tag{D.5}$$

A representative Standard Model example is

$$\frac{d\sigma}{d\Omega}(e^+e^- \to \mu^+\mu^-) = \frac{\alpha^2}{4s}(1 + \cos^2\theta), \qquad \sigma(e^+e^- \to \mu^+\mu^-) = \frac{4\pi\alpha^2}{3s}, \tag{D.6}$$

in the relativistic limit and away from the Z pole.

E Selected observable summary

Observable	Main physical role
Invariant mass m_{ab}	Reveals intermediate resonances, thresholds, and line shapes
Angular distribution	Probes current structure, spin, and interference patterns
Transverse momentum p_T	Natural collider variable relative to the beam axis
Rapidity y / pseudorapidity η	Encodes longitudinal geometry of collider final states
Total rate or yield	Integrates over kinematics and connects directly to luminosity
Branching ratio	Encodes how frequently a parent state chooses a given channel
Resonance width	Connects line shape to lifetime through $\tau = 1/\Gamma$

Two particularly useful definitions are

$$m_{ab}^2 = (p_a + p_b)^2, \qquad \eta = -\ln\tan\frac{\theta}{2}. \tag{E.1}$$

These quantities are part of the standard vocabulary through which theoretical predictions are compared with measured particle-physics data.

F Guided-check summary

Guided checks

- Can you explain the difference between an interaction term, a Feynman rule, a diagram, and an amplitude?
- Can you explain why diagrams are added at amplitude level before squaring?
- Can you state the conceptual difference between dynamics and phase space?
- Can you explain the difference between a partial width, a total width, and a branching ratio?
- Can you explain why a cross section is the natural rate quantity for scattering?
- Can you say why invariant-mass distributions are especially powerful in particle physics?
- Can you explain why Module 6 is a bridge from formal Standard Model structure to experiment rather than a full detector or statistics module?