

The Production of Pairs of Isolated Photons at Higher Orders in QCD

James Christopher Whitehead

Institute for Particle Physics Phenomenology

Department of Physics

Durham University



A thesis submitted to the University of Durham
for the degree of Doctor of Philosophy

2021

The Production of Pairs of Isolated Photons at Higher Orders in QCD

James Christopher Whitehead

Submitted for the degree of Doctor of Philosophy

2021

Abstract: In this thesis, we consider the corrections to the production of a pair of isolated photons at the Large Hadron Collider (LHC) which arise at Next-to-Next-to-Leading-Order (NNLO) in QCD, and Next-to-Leading-Order (NLO) in the electroweak theory. These corrections are calculated through the antenna subtraction formalism, and implemented in the parton-level Monte Carlo program NNLOJET.

This calculation is then applied to a study of the theoretical and phenomenological issues which drive the apparent tension between prior theoretical predictions at this order, and LHC data taken with the ATLAS detector at 8 TeV. In particular, we focus upon the issue of photon isolation, presenting the first calculation of the diphoton process with ‘hybrid isolation’, a compromise between the theoretical and experimental constraints upon predictions and measurements of photonic final-states. We further consider the consequences of another theoretical choice, the renormalisation and factorisation scales at which the calculation is made.

We find that these two theoretical choices act in concert to generate the tension between prediction and data, and show that reasonable alternatives can lead to excellent agreement. We conclude with an application of the same approach to preliminary 13 TeV ATLAS data, again finding excellent agreement.

Contents

Abstract	i
List of Figures	vii
List of Tables	xi
Preface	xix
1 The Standard Model	1
1.1 Symmetries and fields	2
1.1.1 External: spacetime	2
1.1.2 Internal: gauge	4
1.2 Lagrangian dynamics and quantisation	11
1.2.1 Classical field theory	12
1.2.2 Quantum field theory	13
1.2.3 From theory to phenomenology	18
1.3 The Standard Model	20
1.4 Beyond the Standard Model?	30
2 Quantum chromodynamics	35
2.1 The parton model	36
2.1.1 Factorisation	39
2.1.2 DGLAP evolution	40
2.1.3 Testing the parton model	43

2.2	Partonic cross-sections	45
2.3	Divergences in QCD	48
2.3.1	Ultraviolet digergences	49
2.3.2	Infrared divergences	55
2.3.3	Infrared safety	61
2.3.4	PDF renormalisation and mass factorisation	63
2.4	Scale dependence of QCD cross-sections	64
2.4.1	Analytical dependence	65
2.4.2	Estimation of missing higher-orders	66
2.4.3	Scale-setting techniques	68
2.5	Jets and hadrons	70
2.5.1	Jet algorithms	71
2.5.2	Measurement function formalism	74
2.5.3	Jet substructure	75
2.6	The cancellation of IR divergences	76
2.6.1	One-dimensional toy model	77
2.6.2	Application to phase-space integrals	79
3	Antenna subtraction	83
3.1	Antenna functions	84
3.1.1	Crossings	86
3.2	Factorisation of phase-space	88
3.3	Integrated antenna functions	90
3.4	Constructing a subtraction counterterm	91
3.4.1	NLO subtraction counterterms	91
3.4.2	NNLO subtraction counterterms	93
3.5	Application to diphoton production	96
3.5.1	LO	97
3.5.2	NLO	98

3.5.3	NNLO	104
3.6	Implementation of diphoton production in NNLOJET	116
4	Diphoton production: theory	119
4.1	Fragmentation	121
4.2	Isolation	125
4.2.1	Smooth-cone isolation	128
4.2.2	Hybrid isolation	135
4.2.3	Democratic isolation	146
4.2.4	Infrared sensitivity of cone isolation	152
4.3	Cut-dependence	159
5	Diphoton production: phenomenology	163
5.1	Hybrid isolation for phenomenology	165
5.1.1	R_d dependence	165
5.1.2	Comparison of hybrid and smooth-cone distributions	166
5.2	Scale choice	171
5.2.1	Scale choice for photon processes	171
5.2.2	Perturbative convergence	173
5.2.3	Kinematic effects	174
5.2.4	Alternative scale functional forms	179
5.3	Combined effect of isolation and scale variation	182
5.3.1	Comparison to ATLAS data: four-way comparison	183
5.3.2	Comparison to ATLAS data: two-way comparison	185
6	Electroweak corrections	191
6.1	Matrix elements and subtraction terms	193
6.1.1	Real QED radiation	194
6.1.2	Virtual electroweak corrections	195
6.2	Results	197

6.2.1	Impact of electroweak corrections	197
6.2.2	Impact of photon-initiated contributions	199
7	LHC Run 2: outlook at 13 TeV	201
8	Conclusions	211
A	Validation and testing in NNLOJET	213
A.1	Pointwise validation of matrix elements	213
A.2	Validation of subtraction terms	215
A.2.1	Cancellation of real IR divergences	215
A.2.2	Cancellation of virtual ε -poles	215
A.2.3	Finite term consistency checks	217
A.3	Integrated cross-section tests	218
A.3.1	Technical cut dependence	218
A.3.2	Scale variation	219
A.4	Validation against other codes	219
A.4.1	LO and NLO	219
A.4.2	NNLO	221
A.5	Regression tests	224
A.6	Implementation and validation of soft-drop	226
A.6.1	Implementation	226
A.6.2	Validation	228
	Bibliography	229

List of Figures

1.1	The ‘Mexican hat’ potential	23
1.2	Summary of ATLAS cross-section measurements	31
1.3	The evolution of the Unitarity Triangle	33
2.1	Factorisation in QCD	40
2.2	Diagrammatic representation of the DGLAP equations	42
2.3	CT18 NNLO PDFs at 3 and 100 GeV	45
2.4	Experimental verification of the running of the strong coupling	54
2.5	Levels at which a jet algorithm can be applied	72
2.6	Effect of different sequential-clustering jet algorithms	74
3.1	Example Feynman diagrams for diphoton production	97
3.2	Example Feynman diagrams for the double-real contribution	105
3.3	Example Feynman diagrams for the real-virtual contribution	109
3.4	Example Feynman diagrams for the double-virtual contribution	112
4.1	Feynman diagrams for the direct and fragmentation subprocesses	126
4.2	Experimental and unfolded photon isolation at ATLAS	128
4.3	Parameter-dependence of smooth-cone isolation	132
4.4	The matched-hybrid isolation profile function	137
4.5	The variation $\Delta\sigma(R_d) = \sigma_{\text{hybrid}} - \sigma_{\text{smooth}}$ at NLO	139
4.6	Resolution of differences between isolation prescriptions in R and RR kinematics	141

4.7	Isolation cone effects at NLO	142
4.8	Detailed isolation cone effects at NLO	145
4.9	The variation $\Delta\sigma(R_d) = \sigma_{\text{hybrid}} - \sigma_{\text{smooth}}$ at NNLO	146
4.10	Parameter-dependence of soft-drop isolation	150
4.11	Differential comparison of soft-drop and smooth-cone isolation . . .	151
4.12	Discontinuity and Sudakov singularity in $d\sigma/dp_T^{\gamma\gamma}$ at NLO and NNLO	154
4.14	K -factor for diphoton production as a function of the p_T -cuts . . .	160
4.15	The effect of the p_T -cuts on perturbative convergence in the $q\bar{q}$ - and qg -channels.	162
5.1	Fiducial cross-section measured at ATLAS	164
5.3	Direct production with hybrid isolation vs. fragmentation-inclusive with standard isolation	166
5.4	Comparison between matched-hybrid and smooth-cone isolation of $d\sigma/\Delta R_{\gamma\gamma}$ and $d\sigma/M_{\gamma\gamma}$ at NLO and NNLO	167
5.5	R_d -dependence of $d\Delta\sigma/M_{\gamma\gamma}$ and $d\Delta\sigma/p_T^{\gamma\gamma}$ for $\gamma\gamma + j$ production .	169
5.6	Relationship between $p_T^{\gamma\gamma}$ and $\Delta\phi_{\gamma\gamma}$	169
5.7	Comparison of NNLO differential cross-sections with smooth and hybrid isolation	170
5.8	Kinematic regions (x, Q^2) probed according to scale choice	173
5.9	Scale-dependence of the NLO and NNLO cross-sections	175
5.10	Relationship between $M_{\gamma\gamma}$ and $\Delta y_{\gamma\gamma}$	177
5.11	Contributions to the cross-section according to the ratio of scale vari- ables $M_{\gamma\gamma}$ and $\langle p_T^{\gamma\gamma} \rangle$	178
5.12	Scale comparison of differential cross-sections at LO, NLO and NNLO	180
5.13	Four-way scales and isolation comparison at NNLO	183
5.14	Four-way scales and isolation comparison at NNLO, with ATLAS data	186
5.15	Relationship between $ \cos\theta_\eta^* $ and $\Delta\eta_{\gamma\gamma}$	187
5.16	Effect of scale and isolation on ATLAS predictions	188

5.17	Perturbative convergence of distributions measured at ATLAS . . .	190
6.1	NNPDF3.1luxQED NNLO PDFs (inc. f_{γ}^p) at 3 and 100 GeV . . .	194
6.2	Feynman diagrams for NLO EW corrections	196
6.3	Size of NLO electroweak corrections relative to QCD corrections . .	198
6.4	Magnitude of the photon-induced electroweak corrections	200
7.1	NNLO predictions for ATLAS at 13 TeV	203
7.2	Non-analyticity in the NNLO $d\sigma/dM_{\gamma\gamma}$ distribution	204
7.3	NNLO QCD predictions for ATLAS at 13 TeV	206
7.4	NNLO QCD + NLO EW predictions for ATLAS at 13 TeV	207
7.4	NNLO QCD + NLO EW predictions for ATLAS at 13 TeV	208
7.4	NNLO QCD + NLO EW predictions for ATLAS at 13 TeV	209
A.1	Example $4 \rightarrow 2$ double-real spike-tests	216
A.2	NNLOJET scale-variation test	220
A.3	Cut-dependence of the diphoton cross-section within q_T -subtraction	222
A.4	Validation of differential distributions	225

List of Tables

1.1	Standard Model fermion masses	22
3.1	Three-parton antenna functions X_3^0 and X_3^1	85
3.2	Four-parton antenna functions X_4^0	87
3.3	Perturbative order of partonic channels	96
A.1	The parts of the NNLO calculation, and their contribution to other calculations	221
A.2	Cross-section validation at LO, NLO and NNLO	223

Declaration

The work in this thesis is based on research carried out in the Department of Physics at Durham University. No part of this thesis has been submitted elsewhere for any degree or qualification.

Chapters 4 and 5 are based in part on joint research, presented in the publications:

- T. Gehrmann, E.W.N. Glover, A. Huss and J. Whitehead, *Scale and isolation sensitivity of diphoton distributions at the LHC*, *JHEP* **01** (2021) 108 [2009.11310].
- ‘Photon isolation studies’, a contribution to S. Amoroso et al., *Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report*, in *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches*, Mar., 2020 [2003.01700].

In addition, some of the results obtained with the code described in this thesis and presented here have been provided to the ATLAS collaboration and presented in:

ATLAS collaboration, *Measurement of the production cross-section of isolated photon pairs in pp collisions at 13 TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2020-024, CERN, Geneva (July, 2020).

Copyright © 2021 James Christopher Whitehead.

The copyright of this thesis rests with the author. No quotation from it should be published without the author’s prior written consent and information derived from it should be acknowledged.

Acknowledgements

A PhD is a journey, and every journey needs companions. I am especially grateful for mine.

I would like to thank my supervisor Nigel Glover, for his patience, insight and good humour, and fellow PhD students Juan and Duncan for their assistance in establishing OC325 as the IPPP's one-stop-shop for help with all things computing and QCD. I hope I can be as helpful to my successors as they were to me. They, our collaborators Thomas and Alex, and the entire NNLOJET team, are a formidable group of colleagues; I am lucky to have worked with them. I am grateful to the STFC for the funding that made it possible.

The extended ecosystem of the IPPP, from PhD students to staff at all levels, made coming to the office each day a pleasure; special mention to Marian and Alan, the comedy duo always receptive to a technical discussion, and Giuseppe and Gavin, on-hand at home for kitchen physics whenever an interesting problem arose.

Beyond work, my time in Durham would have been unrecognisable without choirs; the glorious SCSCC 'dream team' of Alys and Martha, NSS, DUCS and DCCS have all kept me grounded when not thinking about physics. As have my rotating cast of housemates, especially Antonia, Lindsay, Liz and Max. It's been a pleasure.

And of course, Hannah, my Geordie and guide extraordinaire, who was there before it all began. We've faced a lot, and seen it through. Let's keep exploring. Thanks too to Lynn and Gary, for a warmth and welcome that have proved to this Southerner that It's Great up North.

And finally: my family. You mean the world to me. The last year has been difficult, but things will look up. Thanks for everything.

Truth... is much too complicated to allow anything but approximations.

— J. von Neumann

The Standard Model is so complex that it would be hard to put it on a T-shirt—though not impossible; you'd just have to write kind of small.

— S. Weinberg

Preface

The inauguration of the Large Hadron Collider (LHC) in 2008 heralded the beginning of a new era of precision particle physics. Its successful operation—colliding beams of highly-accelerated protons, measuring the radiation produced, and analysing the torrent of resulting data—is a marvel of modern science and engineering, which has already been rewarded by the confirmed discovery of the Higgs boson in 2012.

The discovery of the Higgs was expected. Its existence was first hypothesised in 1964, and subsequently emerged as the final piece of the electroweak puzzle, completing the Glashow–Weinberg–Salam model of electroweak unification. A ‘no-lose’ theorem guaranteed that the LHC would deliver either a confirmation of the Higgs, or traces of another form of new physics.

Now the Higgs has been discovered, it is unclear whether further new physics remains within the LHC’s reach. We know that the Standard Model cannot be complete, but not what form future additions are likely to take. Our best hope lies in developing our theoretical machinery to a level of precision commensurate with the experimental precision of the LHC era, both to measure the parameters of the known particles in the Standard Model, and to shine light into darkened corners where as-yet-unseen traces of new physics might be hiding.

Since the LHC collides protons, whose constituents belong to the world of Quantum Chromodynamics (QCD), we must focus on precision QCD. The theoretical revolution of the LHC era that has made this feasible is the advent of Next-to-Next-to-Leading Order (NNLO) QCD calculations. These are additional corrections to the perturbative series, of considerable complexity, expected to give closer agreement with the (incalculable) exact prediction implied by QCD.

One reason for this complexity is that QCD, in its massless approximation, contains low-energy, ‘infrared’ divergences in addition to the high-energy ‘ultraviolet’ divergences ubiquitous in quantum field theory. Rendering these divergences compatible with numerical integration is a delicate affair.

‘Antenna subtraction’ uses the universal structure of the infrared divergences to systematically transfer them from numerical integrands to analytic integrals, which have been calculated and can be evaluated as functions of kinematic variables.

In this work, we apply antenna subtraction to the production of two identified, isolated photons, to NNLO in QCD. This is an important process, as a test of perturbative QCD, a background for the measurement of the Higgs, and as a background for possible signals of new physics.

The requirement that the photons be isolated arises on both theoretical and experimental grounds. These different motivations have led to different methods, known as ‘isolation criteria’. The mismatch between these methods leads to an uncertainty in our calculation, which must be properly accounted for if comparisons between theory and data are to be interpreted correctly. We apply our calculation to study this uncertainty both on its own and in combination with another, of distinct origin, the calculational choice of renormalisation and factorisation scale. We then apply these conclusions to the interpretation of data from ATLAS. We further apply the antenna subtraction method to electroweak (EW) corrections, and compare the resulting NNLO QCD + NLO EW prediction to the latest ATLAS 13 TeV data.

The structure of the thesis corresponds to this outline. Concretely, in chapters 1 and 2 we review the basics of the Standard Model and perturbative QCD, and in chapter 3 introduce antenna subtraction and apply it to diphoton production. In chapter 4 we introduce isolation, and study the theoretical consequences of isolation and scale choice for diphoton predictions, which we then apply to ATLAS 8 TeV data in chapter 5. Finally, in chapter 6 we apply antenna subtraction to the calculation of electroweak corrections, which we combine with the NNLO prediction and apply to ATLAS 13 TeV data in chapter 7.

CHAPTER 1

The Standard Model

Particle physics aims to describe the fundamental ‘building blocks’ of the universe: what they are, how they behave, and how we can tell. Its fulfilment of that aim is expected to take the form of a theoretical model, which can—with as little external input as possible—accurately predict and reproduce the results of high-energy experiments testing the interactions of particles at the smallest scales.

Our best candidate for that model is called the ‘Standard Model’ (SM). According to the Standard Model, these building blocks—‘elementary particles’—are best defined not by concrete properties, but by their transformation properties with respect to the symmetries of spacetime, and the symmetries between particles.

In the bulk of this thesis we will work in depth on precise predictions and measurements that aim to test the Standard Model through its consequences for a single process, the production of pairs of isolated photons at hadron colliders. For this to have any hope of describing reality it must be based on a sound theoretical framework.

In this chapter we introduce that framework, building up to the full Standard Model from its underlying symmetries, focusing on the two gauge theories that govern hadrons and photons respectively, quantum chromodynamics (QCD) and quantum electrodynamics (QED). The application of these abstract groups to predictions

for the behaviour of particles through the Lagrangian formalism, quantisation and Feynman diagrams is outlined in section 1.2. Finally in section 1.3 we introduce the Standard Model in full, and in section 1.4 touch upon what may lie beyond.

1.1. Symmetries and fields

1.1.1. External: spacetime

The universe respects both translational and rotational invariance: there is no special location, or special direction, in which the underlying laws are different. Special relativity puts time and space on an equal footing, so we cannot unambiguously distinguish between Lorentz boosts and spatial rotations. Taken together, the group of all these transformations is called the *Poincaré* group. These are the isometries of Minkowski spacetime, which preserve the ‘distance’ between four-vectors¹

$$x \cdot y = x^\mu \eta_{\mu\nu} y^\nu = x_0 y_0 - \mathbf{x} \cdot \mathbf{y}, \quad (1.1.1)$$

where the Minkowski metric $\eta = \text{diag}(+1, -1, -1, -1)$ encodes the hyperbolic geometry of Minkowski spacetime.

To describe the existence and interactions of particles respecting this symmetry, we require a formalism in which their intrinsic properties remain unchanged under it, whilst those related to position or direction do not. This is satisfied if they transform in a representation of the Poincaré group. To satisfy ‘unitarity’ (the conservation of probability), the representation must be unitary; to identify an elementary particle, it should be irreducible.

Thus, the irreducible unitary representations of the Poincaré group dictate what

¹ Here and throughout we use ‘natural’ units, in which two of the three fundamental dimensional constants G , $\hbar$ and c are set to unity:

$$\hbar = c = 1,$$

leaving all dimensionful quantities expressed in a single scale, which we choose to be energy (typically in electronvolts eV). The dimension of a quantity is then reduced to an integer, with mass-dimension d meaning $(\text{mass})^d$.

kinds of elementary particle are possible. They are uniquely classified by just two numbers: ‘mass’ m , and ‘spin’ J .² The apparent absence of particles travelling faster than the speed of light rules out $m < 0$. The remaining representations, and hence particles, with non-negative mass can be classified entirely by mass and spin J , for

$$J = 0, \frac{1}{2}, 1, \frac{3}{2}, \dots \quad (1.1.3)$$

Each such representation contains $2J + 1$ independent states for $m > 0$, and 2 for $m = 0$.

Particles with integer spin are called bosons, and behave differently from particles with half-integer spin, called fermions. All observed elementary fermions are spin- $\frac{1}{2}$. The observed elementary bosons are the Higgs, the only known particle of spin-0, and the photon, the gluon and the W and Z bosons, gauge bosons which carry the electromagnetic, strong and weak forces respectively, of spin-1. These will be introduced properly in section 1.3.

We therefore have all³ the fundamental building blocks of particle theory in the spin-0, spin- $\frac{1}{2}$, and spin-1 irreducible representations of the Poincaré group. The mathematical objects that transform under these representations are called ‘scalars’, ‘spinors’, and ‘vectors’ respectively. An arbitrary Lorentz transformation Λ ,

$$\Lambda = \exp \left(\frac{1}{2} \omega_{\mu\nu} M^{\mu\nu} \right), \quad (1.1.4)$$

for Lie Algebra generators $M^{\mu\nu}$, maps spacetime as

$$x^\mu \mapsto x'^\mu = \Lambda^\mu{}_\nu x^\nu, \quad (1.1.5)$$

² Corresponding to the Casimir invariants $P^\mu P_\mu$ and $W^\mu W_\mu$ respectively, where

$$W_\mu = \frac{1}{2} \varepsilon_{\mu\nu\rho\sigma} J^{\nu\rho} P^\sigma \quad (1.1.2)$$

is the Pauli-Lubanski pseudovector, $J^{\nu\rho}$ is the relativistic angular momentum operator, and P^μ is the four-momentum operator.

³ If gravity can be described through quantum field theory, it will be through a massless spin-2 ‘graviton’. This is touched upon briefly in section 1.4. Massless elementary particles of still higher spin could exist, but conservation laws would prevent them from interacting, so they cannot give rise to long-range forces. Massive composite higher-spin particles do exist, such as the ρ_3 or a_4 light mesons.

and the scalar field $\phi(x)$, spinor field $\psi^\alpha(x)$, and vector field $A^\mu(x)$ each transform under their respective representations as

$$\begin{aligned}\phi(x) &\mapsto \phi'(x) = \phi(\Lambda^{-1}x) \\ \psi^\alpha(x) &\mapsto \psi'^\alpha(x) = S[\Lambda]^\alpha{}_\beta \psi^\beta(\Lambda^{-1}x) \\ A^\mu(x) &\mapsto A'^\mu(x) = \Lambda^\mu{}_\nu A^\nu(\Lambda^{-1}x),\end{aligned}\tag{1.1.6}$$

where

$$S[\Lambda] = \exp\left(\frac{1}{8}\omega_{\mu\nu}[\gamma^\mu, \gamma^\nu]\right),\tag{1.1.7}$$

and the 4×4 ‘Gamma matrices’ $\{\gamma^\mu\}$ are the lowest-dimensional representation of the Clifford algebra defined by the anticommutation relation

$$\{\gamma^\mu, \gamma^\nu\} = 2\eta^{\mu\nu}\mathbf{I}.\tag{1.1.8}$$

1.1.2. Internal: gauge

In contrast to the Poincaré group of spacetime symmetries, gauge symmetries are internal symmetries between particle fields. They arise from the redundancies in our attempt to describe a system with $(2J+1)$ independent physical degrees of freedom with a Lorentz tensor field representation of rank n , if its dimension 4^n is larger.

If $\phi(x)$ transforms under an M -dimensional matrix representation $R(G)$ of an N -dimensional Lie group G ,

$$\phi \rightarrow \phi'(x) = U_R(\boldsymbol{\alpha}) \phi(x),\tag{1.1.9}$$

where $\boldsymbol{\alpha}$ is a vector of N real parameters and

$$U_R(\boldsymbol{\alpha}) = \exp\left[\mathrm{i}g \sum_{a=1}^N \alpha_a R(T^a)\right]\tag{1.1.10}$$

$$= \mathbf{I}_M + \mathrm{i}g \sum_a \alpha_a R(T^a) + \mathcal{O}(g^2).\tag{1.1.11}$$

The ‘generators’ of the group, $\{T^a\}$, form a basis for the Lie algebra of G , $\mathfrak{g}$. For each generator, $R(T^a)$ is an $M \times M$ complex matrix, which must be Hermitian if

the representation is to be unitary, since:

$$\begin{aligned} \mathbf{I}_M &= U_R^\dagger(\boldsymbol{\alpha}) U_R(\boldsymbol{\alpha}) \\ &= \mathbf{I}_M + \mathrm{i}g \sum_a \alpha_a \left[R(T^a) - R(T^a)^\dagger \right] + \mathcal{O}(g^2) \end{aligned} \quad (1.1.12)$$

for all $\boldsymbol{\alpha}$ if and only if $R(T^a)$ is Hermitian. If the determinant of U_R is to be 1, for all $\boldsymbol{\alpha}$, they must also be traceless, since for all complex square matrices A ,

$$\det[\exp A] = \exp[\mathrm{Tr} A]. \quad (1.1.13)$$

This basis of generators may be chosen to be trace-orthogonal in every representation,⁴ normalised so that

$$\mathrm{Tr} \left[R(T^a) R(T^b) \right] =: T_R \delta^{ab}, \quad (1.1.14)$$

where in the fundamental, defining, representation, by convention,

$$T_R = T_F = \frac{1}{2}. \quad (1.1.15)$$

The Lie algebra $\mathfrak{g}$ is characterised by the structure constants $f^{abc} \in \mathbf{R}$, defined through

$$\left[T^a, T^b \right] = \mathrm{i} \sum_{c=1}^N f^{abc} T^c, \quad (1.1.16)$$

since any two Lie algebras with the same structure constants are isomorphic. The generators in the defining representation satisfy

$$T^a T^b = T_F \left[\frac{1}{N} \delta^{ab} \mathbf{I} + \left(d^{abc} + \mathrm{i} f^{abc} \right) T^c \right], \quad (1.1.17)$$

where

$$d^{abc} = \frac{1}{T_F} \mathrm{Tr} \left[\{ T^a, T^b \} T^c \right], \quad f^{abc} = -\frac{\mathrm{i}}{T_F} \mathrm{Tr} \left[[T^a, T^b] T^c \right]. \quad (1.1.18)$$

The structure constants themselves form a representation of dimension N , called the

⁴ This follows, for simple Lie algebras, from the uniqueness of the trace as the only bilinear form $K(a, b)$ that is both symmetric and invariant, satisfying $K([a, b], c) = K(a, [b, c])$ for all $a, b, c \in \mathfrak{g}$. This defines the Cartan–Killing inner product, or form, on $\mathfrak{g}$.

adjoint representation, with components

$$(R_A(T^a))_{bc} = i f^{abc}. \quad (1.1.19)$$

The quadratic Casimir operator C_2 is defined in a general basis and a given representation R as [4, 5]

$$C_2(R) := \frac{1}{T_A} \sum_{a,b,c,d=1}^N f^{acd} f^{bcd} R(T^a) R(T^b), \quad (1.1.20)$$

which commutes with all generators and hence all elements of the algebra. If the generators are orthogonal, eq. (1.1.14) implies that for the adjoint representation,

$$\text{Tr} [R_A(T^a) R_A(T^b)] = - \sum_{c,d=1}^N f^{acd} f^{bdc} = T_A \delta^{ab}, \quad (1.1.21)$$

and so C_2 reduces simply to

$$C_2(R) = \sum_{a=1}^N R(T^a) R(T^a). \quad (1.1.22)$$

Schur's lemma asserts that in an irreducible representation, any such operator must be proportional to the identity, with some constant of proportionality C_R ,

$$C_2(R) =: C_R \mathbf{I}_M. \quad (1.1.23)$$

Taking the trace then implies that

$$N T_R = C_R M. \quad (1.1.24)$$

Such constants will arise regularly in QCD calculations for the fundamental representation, C_F , and the adjoint representation, C_A , for which eq. (1.1.24) gives

$$C_A = T_A. \quad (1.1.25)$$

If α is a global parameter, terms of the form $|\partial_\mu \phi|^2$ and $m^2 |\phi|^2$ are invariant under U , by the unitarity property of eq. (1.1.12). In section 1.2 we will see that these arise as terms governing the kinetic energy and mass of a scalar field ϕ in a Lagrangian density $\mathcal{L}(\phi, \partial_\mu \phi)$. But if we consider a local gauge transformation,

with the parameter α a function of spacetime $\alpha(x)$, the kinetic term ceases to be invariant. To restore the desired invariance we introduce the covariant derivative,

$$D_\mu = \partial_\mu + ig \sum_{a=1}^N V_\mu^a R(T^a) \quad (1.1.26)$$

where V_μ^a are a set of N gauge-vector fields, required to have the transformation law

$$V_\mu^a R(T^a) \rightarrow V_\mu'^a R(T^a) = U_R V_\mu^a R(T^a) U_R^{-1} + \frac{i}{g} \partial_\mu (U_R) U_R^{-1} \quad (1.1.27)$$

using the Einstein summation convention over the index a . Then the covariant derivative of the field, $D_\mu \phi$, has the same transformation properties as ϕ :

$$\begin{aligned} D_\mu \phi &\rightarrow (\partial_\mu + ig V_\mu'^a R(T^a)) \phi' \\ &= (\partial_\mu U_R) \phi + U_R (\partial_\mu \phi) + ig (U_R V_\mu^a R(T^a) U_R^{-1}) (U_R \phi) - (\partial_\mu U_R) U_R^{-1} (U_R \phi) \\ &= U_R (\partial_\mu + ig V_\mu^a R(T^a)) \phi = U_R D_\mu \phi, \end{aligned} \quad (1.1.28)$$

and the modified Lagrangian density $\mathcal{L}(\phi, D_\mu \phi)$ is invariant under gauge transformations, respecting the gauge symmetry.

In order to identify the gauge fields V_μ^a as physical fields, we will need to include a kinetic term in the Lagrangian density, which must include derivatives of the field and be invariant under gauge transformations. This arises naturally from the commutator of covariant derivatives,

$$\begin{aligned} [D_\mu, D_\nu] \phi &= ig [\partial_\mu V_\nu^a - \partial_\nu V_\mu^a + ig f_{abc} V_\mu^b V_\nu^c] R(T^a) \phi \\ &=: ig F_{\mu\nu}^a R(T^a) \phi, \end{aligned} \quad (1.1.29)$$

which defines the field-strength tensor,

$$\mathbf{F}_{\mu\nu} := \sum_{a=1}^N F_{\mu\nu}^a R(T^a). \quad (1.1.30)$$

This tensor represents the physical part of the gauge fields, which cannot be changed by choosing a different gauge. The trace then has the desired properties:

$$\mathcal{L}_{\text{kin}} = -\frac{1}{2} \text{Tr} [\mathbf{F}_{\mu\nu} \mathbf{F}^{\mu\nu}] = -\frac{1}{2} T_R F_{\mu\nu}^a F^{\mu\nu a}, \quad (1.1.31)$$

which will give rise to self-interactions of the gauge field with itself, following from the expansion of eq. (1.1.29). This completes⁵ the so-called Yang–Mills Lagrangian for gauge group $SU(N)$, since gauge invariance implies that there can be no mass term for the gauge fields:⁶

$$\mathcal{L}_{\text{YM}} = -\frac{1}{4}F_{\mu\nu}^a F^{\mu\nu a} + \mathcal{L}(\phi, D_\mu \phi). \quad (1.1.33)$$

This gauge group, and the resulting Lagrangian, is of critical importance to particle physics. The known properties of the Standard Model can be described only with the Yang–Mills theory of the group $SU(N)$ and its abelian sibling, $U(1)$, as will be discussed in further detail in section 1.3.⁷

1.1.2.1. Quantum Electrodynamics

The simplest example of these principles is the abelian gauge group $U(1)$, realised in nature as quantum electrodynamics (QED). The fermionic spinor field transforms as

$$\psi(x) \rightarrow \psi'(x) = \exp(-ieq\alpha(x))\psi, \quad (1.1.34)$$

where q is the eigenvalue of the charge operator for the particle (in units of the electron charge e), and we define its covariant derivative as

$$D_\mu = \partial_\mu + ieqA_\mu(x). \quad (1.1.35)$$

The strength of the electromagnetic interaction is more commonly expressed through the dimensionless ‘fine-structure’ constant α_{em} [7], related (in natural units) to the

⁵ There is a further mass-dimension 4 term, the so-called ‘theta’ term, constructed from the field strength tensor and its dual,

$$\mathcal{L}_\theta \propto \varepsilon_{\mu\nu\rho\sigma} F^{\mu\nu a} F^{\rho\sigma a}. \quad (1.1.32)$$

This can be rewritten as a total derivative of the Chern–Simons current, and so contributes only a boundary term to the action. As we will see later, this implies that it has no effect on the equations of motion, and so we do not consider it further.

⁶ The mechanism by which massive gauge bosons nevertheless arise in nature will be outlined in section 1.3.

⁷ This is slightly surprising, as the Killing–Cartan classification [6] of compact, simple Lie groups implies that theories could also be built from $SO(N)$, $Sp(N)$, and the ‘exceptional’ groups E_6 , E_7 , E_8 , F_4 and G_2 .

elementary charge e by

$$\alpha_{\text{em}} = \frac{e^2}{4\pi} \approx \frac{1}{137}. \quad (1.1.36)$$

The photonic gauge field $A_\mu(x)$ transforms as

$$A_\mu \rightarrow A'_\mu = A_\mu + \frac{1}{e} \partial_\mu \alpha(x), \quad (1.1.37)$$

and the field-strength tensor $F^{\mu\nu}$ reduces in the abelian case to

$$F_{\mu\nu} = \partial_\mu A_\nu - \partial_\nu A_\mu, \quad (1.1.38)$$

so the kinetic term for the photon is

$$-\frac{1}{4} F_{\mu\nu} F^{\mu\nu} \quad (1.1.39)$$

resulting in the QED Lagrangian density

$$\mathcal{L}_{\text{QED}} = -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} + \sum_\psi \bar{\psi} (\mathrm{i}\gamma \cdot D - m_\psi) \psi. \quad (1.1.40)$$

1.1.2.2. Quantum Chromodynamics

Quantum Chromodynamics (QCD) is based on the gauge group $\text{SU}(3)$. The triplets, transforming in the fundamental representation, are called ‘quarks’, and the octets, transforming in the adjoint representation, are called ‘gluons’.

In accordance with eq. (1.1.33) its classical Lagrangian density is

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} F_{\mu\nu}^a F^{a\mu\nu} + \sum_{\text{quarks } q} \bar{\psi}_q^i (\mathrm{i}\gamma \cdot D_{ij} - m_q \delta_{ij}) \psi_q^j, \quad (1.1.41)$$

where i and j are colour indices (in the fundamental representation), D_{ij}^μ is the QCD covariant derivative,

$$[D_\mu]_{ij} = \delta_{ij} \partial_\mu - \mathrm{i} g_s T_{ij}^a G_\mu^a(x), \quad (1.1.42)$$

and $\{T^a\}_{a=1}^8$ are the standard Gell-Man generators of $\text{SU}(3)$,

$$T^a = \frac{1}{2} \lambda^a, \quad (1.1.43)$$

normalised to satisfy eq. (1.1.15).

The generators $\{T^a\}_{a=1}^{N^2-1}$ of $\text{SU}(N)$ form a basis for $\mathfrak{su}(N)$, which can be extended to a basis for the Hermitian matrices by the inclusion of a single additional linearly-independent matrix, such as the identity matrix. Expanding in this basis and using the tracelessness and trace-orthogonality of the generators to extract the coefficients gives, for all Hermitian matrices M , the ‘completeness relation’

$$M = \text{I}_N \left(\frac{1}{N} \text{Tr } M \right) + \sum_{a=1}^{N^2-1} T^a \left(\frac{1}{T_F} \text{Tr } [MT^a] \right). \quad (1.1.44)$$

Expressed in component notation, this implies that

$$M_{lk} \left[\delta_{il} \delta_{jk} - \frac{1}{N} \delta_{ij} \delta_{kl} - \frac{1}{T_F} \sum_{a=1}^{N^2-1} T_{ij}^a T_{kl}^a \right] = 0, \quad (1.1.45)$$

for all M_{lk} , and hence the ‘Fierz identity’

$$\sum_{a=1}^{N^2-1} T_{ij}^a T_{kl}^a = T_F \left[\delta_{il} \delta_{jk} - \frac{1}{N} \delta_{ij} \delta_{kl} \right]. \quad (1.1.46)$$

This fixes T_A in terms of T_F , as it implies

$$(N^2 - 1) T_A = \sum_{a,b,c=1}^{N^2-1} f^{abc} f^{abc} \quad (1.1.47)$$

$$= \sum_{a,b=1}^{N^2-1} -\frac{1}{T_F} \text{Tr} \left[[T^a, T^b] [T^a, T^b] \right] = 2T_F N(N^2 - 1), \quad (1.1.48)$$

and so

$$T_A \equiv C_A = 2T_F N. \quad (1.1.49)$$

With the $T_F = \frac{1}{2}$ convention of eq. (1.1.15), and using eq. (1.1.24) the Casimir invariants of eq. (1.1.23) for the fundamental and adjoint representations are therefore

$$C_F = \frac{N^2 - 1}{2N} \quad C_A = N. \quad (1.1.50)$$

For QCD, with $N_c = 3$, this gives $C_F = \frac{4}{3}$ and $C_A = 3$. These factors will arise frequently in calculations.

The constant g_s in eq. (1.1.41) is called the ‘strong coupling’, more commonly

expressed in terms of the dimensionless quantity α_s through the relationship

$$\alpha_s = \frac{g_s^2}{4\pi}. \quad (1.1.51)$$

In practice, as a result of renormalisation, this coupling is not in fact constant, but a function $\alpha_s(\mu)$ of an energy scale μ . Its variation as a function of the scale is called the ‘running of the coupling’ and is an important feature of perturbative QCD, described further in chapter 2.

The fact that the mass matrix in eq. (1.1.41) is diagonal corresponds to a conventional choice of quark fields as the eigenstates of the free Hamiltonian. In the context of the full Standard Model, these eigenstates differ from those of the electroweak interactions, which gives rise to the Cabibbo–Kobayashi–Maskawa (CKM) matrix parametrising this change of basis. This will be discussed further in section 1.3.

In practice however, in chapter 2 we will work in the effective theory of massless QCD, in the high-energy limit in which the lightest n_f quark masses (listed in table 1.1) become negligible, and so these ‘light’ quarks are assumed to be massless. This is accurate up to corrections suppressed by powers of m_q/E , provided that precautions are taken to keep the limit well-defined and free from singularities in the infrared or massless limits. These will be discussed further in section 2.3.3.

1.2. Lagrangian dynamics and quantisation

Lagrangian densities such as those described in section 1.1 describe the physical behaviour of particles through the framework of quantum field theory. The Lagrangian density itself is sufficient to describe the theory classically, through the principle of least action.

1.2.1. Classical field theory

The *principle of least action* asserts that the path taken by a system is that in which the classical action,

$$S[\{\phi_a\}] = \int_V d^4x \mathcal{L}(\{\phi_a\}; \{\partial_\mu \phi_a\}) \quad (1.2.1)$$

is extremised, $\delta S = 0$, requiring that the Euler–Lagrange equations,

$$\partial_\mu \left(\frac{\partial \mathcal{L}}{\partial (\partial_\mu \phi_a)} \right) = \frac{\partial \mathcal{L}}{\partial \phi_a}, \quad (1.2.2)$$

are satisfied. This gives the equations of motion for the evolution of the fields. As a result, any two Lagrangian densities $\mathcal{L}$ and $\mathcal{L}'$ differing by a total derivative must lead to the same equations of motion, since through the divergence theorem the extra term contributes to S only on the bounding surface ∂V , not in the bulk. This leads to Noether’s theorem, in which symmetries of the Lagrangian lead to conservation laws.⁸

When applied to the Dirac Lagrangian for a free fermion field,

$$\mathcal{L}_D = \bar{\psi} (i\gamma \cdot \partial - mI) \psi, \quad (1.2.3)$$

which can be identified as part of both the QED and the QCD Lagrangians, the Euler–Lagrange equations give the Dirac equation,

$$(i\gamma \cdot \partial - mI) \psi(x) = 0. \quad (1.2.4)$$

Similarly, the source-free Lagrangian density for classical electromagnetism, given by the kinetic term for the QED photon of eq. (1.1.39), gives Maxwell’s equations for classical electromagnetism,

$$\partial_\mu F^{\mu\nu} = 0. \quad (1.2.5)$$

⁸ These conservation laws include the deep result that energy and momentum conservation are a consequence of translation invariance, and angular momentum and centre-of-momentum conservation are a consequence of Lorentz invariance.

1.2.2. Quantum field theory

The formulation of the corresponding quantum theory requires the additional ingredients of quantisation, regularisation, and renormalisation. Non-abelian gauge theories require the further introduction of fictitious ‘ghost’ fields, additional terms in the Lagrangian density constructed to eliminate the redundant degrees of freedom in our description of gauge symmetries, as mentioned in section 1.1.2.

Quantum field theory is rich and we will not attempt to review it in its full detail; detailed reviews are available in [8–10]. Instead we will sketch the important features and summarise the conclusions that lead to the possibility of computing physical quantities through perturbation theory.

1.2.2.1. Canonical quantisation

Fields of different spins (scalars, fermions and vector bosons) are each quantised differently, but all have common features. The classical fields $\phi(\mathbf{x}, t)$ and conjugate momenta $\pi(\mathbf{x}, t)$ are replaced by field operators $\hat{\phi}(\mathbf{x}, t)$ and $\hat{\pi}(\mathbf{x}, t) = i\hat{\phi}^\dagger(\mathbf{x}, t)$, which are required to satisfy equal-time commutation or anti-commutation relations according to whether the corresponding particle should obey Bose-Einstein or Fermi-Dirac statistics. The dynamics of the field operators is then given by Heisenberg’s operator equation of motion,

$$\dot{\hat{\phi}}(\mathbf{x}, t) = -i[\hat{\phi}, \hat{H}], \quad \dot{\hat{\pi}}(\mathbf{x}, t) = -i[\hat{\pi}, \hat{H}] \quad (1.2.6)$$

which reproduces for the field operators the Euler–Lagrange equations of motion satisfied classically by the fields.

Expanding the field operators as a superposition of plane-waves gives a general solution to the operator equation of motion, subject to an (anti-)commutation constraint on the expansion coefficients. Defining the ‘normal ordering’ of operators to remove the divergent zero-point field energy, we can exploit these (an-

ti-)commutation relations to define a positive-definite Hamiltonian.⁹ The resulting Hamiltonian resembles that of the quantum harmonic oscillator for each value of the three-momentum $\mathbf{p}$ and admits a similar interpretation, of creation and annihilation operators that act successively on the vacuum-state $|0\rangle$ to form a multi-particle state with the specified momenta. Such states form a basis for the Hilbert space of multiple-particle states, called the Fock space.

1.2.2.2. Interactions

‘Free’ Lagrangians are bilinear (or sesquilinear) in the fields, and the coefficient of the quadratic term defines the mass of the field. Higher-order terms than this define ‘interactions’. These may be self-interactions, arising from higher powers of a single field, or interactions between different fields arising from products of their field operators. This seemingly-modest complication is enough to prevent the dynamics of interacting theories from being exactly solvable, as they are for free fields.

However, if the strength of these interactions is sufficiently small, the dynamics can be solved approximately as a perturbation of the free theory, decomposing the Hamiltonian (density) into a free and an interacting part as

$$\hat{H} = \hat{H}_0 + \lambda \hat{H}_1. \quad (1.2.7)$$

The aim of perturbation theory is to derive a formal power series for each observable, in the parameter $\lambda \ll 1$ defining the strength of the interaction. As we compute successive terms, the approximation of the perturbative solution to the exact one should improve,¹⁰ allowing successively-refined predictions for comparison to experiment.

⁹ If commutation relations (implying Bose–Einstein statistics) are used for the Dirac field instead of anti-commutation relations (for Fermi–Dirac) this is impossible, as the spectrum of the resulting Hamiltonian is not bounded from below.

¹⁰ In fact, this is an asymptotic series with vanishing radius of convergence, so if we include enough terms the series will eventually diverge. This is possible because asymptotic series do not uniquely define the function they approximate, being insensitive to contributions whose derivatives all vanish at $\lambda = 0$ such as $e^{-1/a\lambda}$. Such contributions are called ‘instantons’ and are non-perturbative, so we do not consider them further. We can deduce, however, that these missing contributions do not become significant until the order $k \approx \frac{1}{\lambda}$, which for the theories we shall consider is far higher than we can calculate. We therefore expect to see convergent behaviour in the terms which we can.

This power series arises ultimately from the operator equation of motion for the time-evolution operator in the interaction picture,

$$\partial_t \hat{U}(t, t_0) = -i\lambda \hat{H}_1^I(t) \hat{U}(t, t_0), \quad (1.2.8)$$

which can be solved by the iterated integral

$$\hat{U}(t, t_0) = \mathbf{I} + \sum_{n=1}^{\infty} (-i\lambda)^n \int_{t_0}^t dt_1 \cdots \int_{t_0}^{t_{n-1}} dt_n \prod_{m=1}^n \hat{H}_1^I(t_m). \quad (1.2.9)$$

By reorganising the simplicial integration region this can be rewritten as Dyson's formula, where all products of the Hamiltonian are time-ordered,

$$\hat{U}(t, t_0) = \mathbf{I} + \sum_{n=1}^{\infty} \frac{(-i\lambda)^n}{n!} \int_{t_0}^t dt_1 \cdots \int_{t_0}^t dt_n \mathbf{T} [\hat{H}_1^I(t_1) \cdots \hat{H}_1^I(t_n)]. \quad (1.2.10)$$

Through Wick's theorem, these time-ordered products can be written as normal-ordered products of contracted fields. These can be visualised intuitively as (position-space) Feynman diagrams, with each element of a diagram corresponding to a factor appearing in the integral. The simple (plane-wave) x -dependence allows the d^4x integrations from the interaction Hamiltonian densities to be integrated over, giving equivalent momentum-space Feynman rules. These are the diagrams we use in practice. We will not list the rules here, but they can be found in any introduction to QFT and the Standard Model (for example, [9–13] or in full, listed for all possible sign conventions, in [14]).

Although Feynman diagrams are useful calculational and conceptual tools, pursuing this approach at higher orders in λ leads to computational difficulties. This poses little trouble for QED, where the small coupling constant gives fast convergence, but becomes problematic for QCD, whose larger coupling constant necessitates higher-order expansions in the perturbative series. In particular, the difficulty of calculating these diagrams increases with their complexity, and the number of diagrams that must be calculated increases factorially, whilst the eventual sum shows remarkable simplifications [15, 16].

A particularly fruitful alternative is the spinor-helicity formalism [17], which

exploits the fact that lightlike momentum vectors p^μ can be factorised as an outer product of spinors (in the bispinor representation) to work in a basis of two-component spinors of definite helicity,

$$p^{\alpha\dot{\beta}} := p^\mu \sigma_\mu^{\alpha\dot{\beta}} = \begin{pmatrix} p^0 - p^3 & -p^1 + ip^2 \\ -p^1 - ip^2 & p^0 + p^3 \end{pmatrix} \quad (1.2.11)$$

$$\equiv |p\rangle^\alpha [p]^{\dot{\beta}}, \quad (1.2.12)$$

where the Pauli matrix four-vector is given by embedding the three Pauli matrices σ^a into Minkowski space,

$$(\sigma^\mu) = \begin{pmatrix} \mathbf{I}_2 & \sigma^1 & \sigma^2 & \sigma^3 \end{pmatrix}, \quad (1.2.13)$$

and the spinors are given by

$$|p\rangle^\alpha := \frac{t}{\sqrt{p^0 - p^3}} \begin{pmatrix} p^0 - p^3 \\ -p^1 - ip^2 \end{pmatrix} \quad (1.2.14)$$

$$[p]^{\dot{\beta}} := \frac{t^{-1}}{\sqrt{p^0 - p^3}} \begin{pmatrix} p^0 - p^3 & -p^1 + ip^2 \end{pmatrix}. \quad (1.2.15)$$

This dramatically simplifies the number of diagrams that must be calculated because many helicity configurations trivially vanish or are related, but in conventional Feynman techniques full generality is maintained only to simplify in a final sum over the spins of external particles. For example, for a seven-gluon tree-level amplitude, there are around 2500 Feynman diagram, but only nine independent helicity configurations; for 10 gluons there are over 10 million.¹¹

These techniques will be expanded on briefly in section 2.3 in the context of colour-ordered QCD amplitudes. Here we remark on a remarkable consequence of the spinor-helicity formalism that arises from group-theory alone. The four-momentum of a massless particle can be parametrised using spherical coordinates

¹¹These can be counted with a simple recurrence relation, or by looking up OEIS sequence A268163.

as

$$p^\mu = E \begin{pmatrix} 1 & \sin \theta \cos \phi & \sin \theta \sin \phi & \cos \theta \end{pmatrix}, \quad (1.2.16)$$

giving spinor variables

$$|p\rangle^\alpha = t\sqrt{2E} \begin{pmatrix} \sin \frac{\theta}{2} \\ -e^{i\phi} \cos \frac{\theta}{2} \end{pmatrix} \quad (1.2.17)$$

$$[p]_{\dot{\beta}} = t^{-1}\sqrt{2E} \begin{pmatrix} \sin \frac{\theta}{2} & -e^{-i\phi} \cos \frac{\theta}{2} \end{pmatrix}. \quad (1.2.18)$$

In the momentum-aligned frame, given by $\theta = 0$,

$$p^\mu = \begin{pmatrix} E & 0 & 0 & E \end{pmatrix}, \quad (1.2.19)$$

and so is manifestly invariant under the isometry group of the transverse xy -plane, $E(2)$. This corresponds to Wigner's 'little group' [18], the subgroup of Poincaré transformations that leave the particle's momentum unchanged. In this frame, this includes the subspace

$$SO(2) \cong U(1) \quad (1.2.20)$$

of rotations about the z -axis, parametrised in eq. (1.2.17) by the rotation angle ϕ .

In the spinor-helicity variables, $\theta = 0$ implies

$$|p\rangle^\alpha = -te^{i\phi}\sqrt{2E} \begin{pmatrix} 0 \\ 1 \end{pmatrix} \quad [p]_{\dot{\beta}} = -t^{-1}e^{-i\phi}\sqrt{2E} \begin{pmatrix} 0 & 1 \end{pmatrix}, \quad (1.2.21)$$

and so any expression containing products of $|p\rangle$ and $[p]$, and invariant under the little group on physical grounds, must be invariant under the scaling

$$[p]_a \rightarrow t [p]_a, \quad |p\rangle^{\dot{a}} \rightarrow t^{-1} |p\rangle^{\dot{a}}, \quad (1.2.22)$$

where the choice of t can be seen to be equivalent to the action of the little-group in the momentum-aligned frame of p^μ . Although for real p^μ , t and $e^{i\phi}$ are both constrained to be a complex phase by the requirement

$$(|p\rangle^\alpha)^* \equiv [p]_{\dot{\alpha}} \quad ([p]_{\dot{\beta}})^* \equiv \langle p|_\beta, \quad (1.2.23)$$

imposed on solutions to the Dirac equation by conjugation, in complex kinematics this restriction is lifted and $t, e^{i\phi} \in \mathbb{C}$.

This principle, together with dimensional analysis and the locality of the Lagrangian, is sufficient to uniquely determine the massless 3-particle amplitudes [17]. Analytic recursion relations then allow more complicated amplitudes to be built up from simpler ones [16, 19]. These approaches have led to a revolution in the calculation of the amplitudes necessary for higher-order QCD calculations.

1.2.3. From theory to phenomenology

The formalism of the preceding sections was necessarily abstract. Its validity as a theory for describing the real world can only be tested through falsifiable predictions that can be compared to the outcome of experiments.

A natural measurement to consider is the ‘cross-sectional area’, or ‘cross section’ σ , inspired by Rutherford’s pioneering calculation of the size of a gold nucleus. For scattering of a beam of particles off a single nucleus,

$$\sigma = \frac{\text{number of scattered particles}}{\text{time} \cdot \text{beam number density} \cdot \text{beam velocity}} = \frac{N}{T\Phi} \quad (1.2.24)$$

where Φ is the incoming flux. The number of particles that are at all deflected by short-distance interactions between the incoming and target particle is then proportional to the cross-section, with the constant of proportionality specific to each experiment. For further detail about the shape of the object or potential we might consider the *differential* cross-section, representing the cross-section scattered, per unit solid angle, into each direction (θ, ϕ) ,

$$\frac{d\sigma}{d\Omega}, \quad (1.2.25)$$

from which the total cross-section is recovered through the integral over the solid angle,

$$\sigma = \int_0^{2\pi} d\phi \int_0^\pi d\theta \sin \theta \frac{d\sigma}{d\Omega}. \quad (1.2.26)$$

In the quantum theory particles are no longer point-like objects of definite extent, and so the particles interact only with some probability. We therefore generalise the concept of the cross-section to take into account the probability of the interaction, as

$$d\sigma = \frac{1}{T} \frac{1}{\Phi} dP. \quad (1.2.27)$$

The differential probability dP is precisely what emerges from quantum field theory. We assume that asymptotic states are approximately free, and so the initial and final states $|i\rangle$ and $|f\rangle$ are eigenstates of the free theory. The S -matrix, or scattering matrix, between the two states is then given by

$$\langle f|\hat{S}|i\rangle := \lim_{t_2 \rightarrow \infty} \lim_{t_1 \rightarrow -\infty} \langle f; t_2 | i; t_1 \rangle_S \quad (1.2.28)$$

where the final S denotes the Schrödinger picture, as opposed to the Heisenberg picture on the left-hand side, and so

$$\hat{S} = \lim_{t_2 \rightarrow \infty} \lim_{t_1 \rightarrow -\infty} \hat{U}(t_2, t_1). \quad (1.2.29)$$

In terms of the S -matrix, the differential scattering probability is

$$dP = \frac{|\langle f|\hat{S}|i\rangle|^2}{|\langle i|i\rangle||\langle f|f\rangle|} d\Pi_n, \quad (1.2.30)$$

where $d\Pi_n$ is the phase space measure for an n -particle final state,¹²

$$d\Pi_n = \prod_{i=1}^n V \frac{d^3 p_i}{(2\pi)^3}. \quad (1.2.31)$$

Clearly in a free theory, the S -matrix is the identity matrix. We can therefore isolate the interaction from the free theory by defining the transfer matrix T as

$$i\hat{T} = (\hat{S} - \hat{I}). \quad (1.2.32)$$

¹²The volume factor V is necessary for the proper normalisation of the phase-space, and also appears in the definition of the flux and normalised states. Between these factors of different origins, the dependence of $d\sigma$ on V cancels, allowing the $V \rightarrow \infty$ limit to be taken without difficulty. Details of this calculation can be found in [10].

To make explicit the necessary momentum conservation between initial-state and final-state particles, we factorise T further:

$$i\hat{T} = (2\pi)^4 \delta^{(4)} \left(\sum_i p_i \right) (i\hat{M}). \quad (1.2.33)$$

The problem is therefore essentially reduced to calculating matrix-elements $\langle f | \hat{M} | i \rangle$, either through Feynman rules or other approaches. For $|i\rangle \neq |f\rangle$, this is the only contribution to the S -matrix, and so we typically talk about the squared matrix-element corresponding to identified initial and final states,

$$|M_{i \rightarrow f}|^2 = |\langle f | \hat{M} | i \rangle|^2. \quad (1.2.34)$$

The validity of the underlying theory is then tested experimentally through measurements of σ , the cross-section, and its differential counterparts $d\sigma/dX$, where some dimensions of the final-state phase-space remain to be integrated over, and residual dependence on X remains.

At modern high-energy particle colliders, this is complicated by the fact that two beams collide, rather than a single beam with a fixed target. At the Tevatron and Large Hadron Collider, it is complicated still further, as the accelerated beams are not beams of elementary particles, but beams of composite hadrons (specifically, protons p and anti-protons $\bar{p}$). This makes it necessary to formulate a theory of how the parts (imaginatively named partons) relate to the composite whole. For perturbative QCD, this takes the form of the ‘parton model’, and a universal factorisation ansatz that will be discussed further in chapter 2.

1.3. The Standard Model

The Standard Model Lagrangian density can be written, succinctly, as

$$\begin{aligned} \mathcal{L}_{\text{SM}} = & -\frac{1}{4} F_{\mu\nu} F^{\mu\nu} \\ & + i\bar{\psi} (\gamma \cdot D) \psi \end{aligned}$$

$$\begin{aligned}
& + y_{ij} \bar{\psi}_i H \psi_j + \text{h.c.} \\
& + \left| D_\mu H \right|^2 - V(H),
\end{aligned} \tag{1.3.1}$$

where each line respectively describes gauge fields and their self-interactions, fermions and their interactions with gauge fields, the interaction of fermion fields with the Higgs boson, and the Higgs and its self-interactions.

It is a gauge field theory with symmetry group

$$\text{SU}_c(3) \times \text{SU}_L(2) \times \text{U}_Y(1), \tag{1.3.2}$$

and fields of spin 0, $1/2$ and 1.

The first subscript, c , refers to *colour*, and to quantum chromodynamics. Fermions which transform trivially under $\text{SU}_c(3)$ are called leptons (ℓ). Each type of fermion, for both quarks and leptons, comes in three generations, of increasing mass and given distinctive names, as shown in table 1.1. For example, the ‘up’ and ‘down’ quarks of the second generation are called ‘charm’ c and ‘strange’ s quarks respectively, and those of the third generation are called ‘top’ t and ‘bottom’ b .

The second subscript, L , refers to left-handedness: only left-handed fermions (and right-handed antifermions) carry this ‘weak isospin’ quantum number; it is otherwise zero. The generators of $\text{SU}_L(2)$, which we shall denote t^a , are the canonically-normalised Pauli matrices,

$$t^a = \frac{1}{2} \sigma^a, \tag{1.3.3}$$

and so satisfy

$$t^a t^b = \frac{1}{4} \delta^{ab} \text{I}_2 + \frac{i}{2} \epsilon^{abc} t^c. \tag{1.3.4}$$

The final subscript Y refers to weak hypercharge and serves as a mnemonic to distinguish $\text{U}_Y(1)$ from the isomorphic group associated with electromagnetic charge and QED, $\text{U}_{\text{em}}(1)$. The subgroup

$$\text{SU}_L(2) \times \text{U}_Y(1) \tag{1.3.5}$$

Generation	Quark	Mass (MeV)	Lepton	Mass (MeV)
I	u	2.16	e	0.510 998 946
	d	4.67	ν_e	< 0.0011
II	c	1270	μ	105.658 374 5
	s	93	ν_μ	< 0.19
III	t	172 760	τ	1776.86
	b	4180	ν_τ	< 18.2

Table 1.1: Central values for the Standard Model quark and charged lepton masses, and upper bounds for the neutrino masses, as given by the Particle Data Group [22].

is called the electroweak sector, with $SU_L(2)$ gauge bosons W_μ^a , hypercharge gauge boson B_μ , and the Lagrangian density

$$\begin{aligned} \mathcal{L}_{\text{EW}} = & -\frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \\ & + |D_\mu H|^2 - V(H), \end{aligned} \quad (1.3.6)$$

where $W_{\mu\nu}^a$ and $B_{\mu\nu}$ are the field strengths of W_μ^a and B_μ respectively, as defined in eq. (1.1.29), and the form of the covariant derivative follows from eq. (1.1.26),

$$D_\mu H := \left(\partial_\mu - ig W_\mu^a t^a - \frac{i}{2} g' B_\mu \right) H, \quad (1.3.7)$$

where g and g' are the coupling constants for $SU_L(2)$ and $U_Y(1)$ respectively. The electroweak sector undergoes ‘spontaneous symmetry breaking’, to give the physical bosons that mediate the weak interaction, $W^\pm$ and Z , and the electromagnetic photon field γ we know from QED. This Spontaneous Symmetry Breaking (SSB), through the Higgs(–Englert–Brout) mechanism [20,21], is the means by which matter fields acquire mass despite the absence of an explicit mass term in eq. (1.3.1), which would break gauge symmetry.

The Higgs potential in the final line of eq. (1.3.1) has the form

$$\begin{aligned} V(H) = & -\mu^2 (H^\dagger H) + \lambda (H^\dagger H)^2, & \mu^2, \lambda > 0 \\ = & -\mu^2 |H|^2 + \lambda |H|^4, \end{aligned} \quad (1.3.8)$$




Figure 1.1: The Higgs (‘Mexican hat’) potential $V(\phi)$ for complex scalar field ϕ which leads to spontaneous symmetry breaking. The vacuum state lies in the valley around the brim of the hat.

often called the ‘Mexican-hat potential’, and illustrated in fig. 1.1. The ground state at which the potential is minimised is given for

$$|H| = \sqrt{\frac{\mu^2}{2\lambda}} =: \frac{v}{\sqrt{2}}, \quad (1.3.9)$$

and is degenerate, with a residual $U_{\text{em}}(1)$ symmetry

$$V\left(e^{i\alpha} \frac{v}{\sqrt{2}}\right) = V\left(\frac{v}{\sqrt{2}}\right) = -\frac{\mu^4}{4\lambda} \quad \forall \alpha \in [0, 2\pi]. \quad (1.3.10)$$

We can expand the H field about a minimum, without loss of generality

$$H(x) = \frac{1}{\sqrt{2}} e^{it^a \xi^a(x)/v} \begin{pmatrix} 0 \\ v + h(x) \end{pmatrix}, \quad (1.3.11)$$

choosing a direction for the vacuum expectation value of the field,

$$\langle 0|H|0\rangle = \frac{v}{\sqrt{2}} \begin{pmatrix} 0 \\ 1 \end{pmatrix}. \quad (1.3.12)$$

Substituting this expression for H into the SM Lagrangian density gives terms corresponding to Feynman rules for the physical Higgs field $h(x)$, its mass term (with mass $m_h = \sqrt{2}\mu$), mass terms and Feynman rules for the electroweak gauge

bosons W^+ , W^- and Z , and Feynman rules for the massless QED photon γ . For example, direct substitution into the Higgs potential $V(H)$ gives cubic and quartic self-interaction terms for the Higgs,

$$\mathcal{L}_{\text{EW}} \ni -V(H) \ni -\frac{\lambda}{4} \left(h(x)^4 + 4h(x)^3 v + \dots \right). \quad (1.3.13)$$

The kinetic term for the Higgs,

$$|D_\mu H|^2 \ni \frac{1}{8} g^2 v^2 \left[W_\mu^1 W^{1\mu} + W_\mu^2 W^{2\mu} + \left(W_\mu^3 - \frac{g'}{g} B_\mu \right)^2 \right] \quad (1.3.14)$$

gives the mass terms for three massive gauge bosons. The third boson field, Z , is defined by

$$Z_\mu = \frac{e}{g} W_\mu^3 - \frac{e}{g} B_\mu, \quad (1.3.15)$$

so the boson masses can be read off as

$$M_W = \frac{v}{2} g \qquad M_Z = \frac{v}{2} \frac{gg'}{e}, \quad (1.3.16)$$

whilst the orthogonal combination gives the other neutral boson, the QED photon,

$$A_\mu = \frac{e}{g} W_\mu^3 + \frac{e}{g'} B_\mu, \quad (1.3.17)$$

which has no mass term and is consequently massless. The terms of the Lagrangian corresponding to eq. (1.3.14), with $h(x)$ in place of v , give interactions of the gauge boson fields with the physical Higgs boson.

If the electroweak gauge symmetry were a global symmetry, the fields $\xi^a(x)$ corresponding to the generators of the broken symmetry t^a in eq. (1.3.11) would each become massless Nambu-Goldstone bosons, as a consequence of Goldstone's theorem [23]. Instead they must be unphysical, as they can be removed through the gauge transformation

$$H(x) \rightarrow H'(x) = e^{-it^a \xi^a(x)/v} H(x). \quad (1.3.18)$$

These three unphysical degrees of freedom provide the additional degree of freedom

required to move each of the $W^\pm$ and Z fields from the massless spin-1 representation of the Poincaré group, with two polarisations, to the massive spin-1 representation with three, as discussed in section 1.1.1.

The transformations of the gauge boson fields in eqs. (1.3.15) and (1.3.17) to diagonalise the mass-matrix amount to a rotation, by the ‘Weinberg’ (or ‘weak-mixing’) angle θ_w , with

$$g \sin \theta_w = e = g' \cos \theta_w. \quad (1.3.19)$$

Applying this rotation to the kinetic terms for $W_{\mu\nu}^3$ and $B^{\mu\nu}$ in eq. (1.3.6) gives

$$\mathcal{L}_{\text{EW,kin}} = -\frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \quad (1.3.20)$$

$$\ni -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} - \frac{1}{4}Z_{\mu\nu}Z^{\mu\nu}, \quad (1.3.21)$$

where

$$Z_{\mu\nu} := \partial_\mu Z_\nu - \partial_\nu Z_\mu \quad (1.3.22)$$

and we can identify the first term as the kinetic term for the photon from the QED Lagrangian eq. (1.1.39). To identify the physical $W^\pm$ bosons, we look for eigenstates of the QED charge operator,

$$\hat{Q} = t^3 + \frac{1}{2}\hat{Y}. \quad (1.3.23)$$

Since eq. (1.3.4) implies that

$$[t^3, t^1 \pm i t^2] = i(t^2 \mp i t^1) = \pm(t^1 \pm i t^2), \quad (1.3.24)$$

the combinations

$$W_\mu^\pm = \frac{1}{\sqrt{2}}(W_\mu^1 \pm i W_\mu^2) \quad (1.3.25)$$

have charges of ± 1 . Re-expressed in terms of the fields for the physical gauge bosons

A_μ , Z_μ and $W_\mu^\pm$, and applying eq. (1.1.29), the kinetic terms of eq. (1.3.6) become

$$\mathcal{L}_{\text{EW,kin}} = -\frac{1}{4}W_{\mu\nu}^a W^{a\mu\nu} - \frac{1}{4}B_{\mu\nu}B^{\mu\nu} \quad (1.3.26)$$

$$\begin{aligned} &= -\frac{1}{4} \left[2W_{\mu\nu}^+ W^{-\mu\nu} + Z_{\mu\nu}Z^{\mu\nu} + F_{\mu\nu}F^{\mu\nu} \right. \\ &\quad + g \left[(W_{\mu\nu}^+ + W_{\mu\nu}^-) (W^{+\mu} - W^{-\mu}) (Z^\nu \cos \theta_w + A^\nu \sin \theta_w) + \dots \right] \\ &\quad \left. - g^2 \left[4W_\mu^+ W^{-\mu} Z_\nu Z^\nu \cos^2 \theta_w - 8W_\mu^+ W^{-\mu} Z_\nu A^\nu \sin \theta_w \cos \theta_w + \dots \right] \right] \end{aligned} \quad (1.3.27)$$

where here

$$W_{\mu\nu}^\pm := \partial_\mu W_\nu^\pm - \partial_\nu W_\mu^\pm, \quad (1.3.28)$$

and where for brevity additional terms carrying the same gauge-theory factor have been omitted (the full expansion can be found in [24]). The latter two lines give terms corresponding to interactions of three gauge bosons, containing a single derivative of a field from the abelian part of one covariant derivative, such as

$$\mathcal{L}_{\text{EW,kin}} \ni -ig \cos \theta_w W_{\mu\nu}^+ W^{-\mu} Z^\nu, \quad (1.3.29)$$

and quartic interaction terms containing no field derivatives from the product of the non-abelian parts of both covariant derivatives, such as

$$\mathcal{L}_{\text{EW,kin}} \ni -g^2 \cos^2 \theta_w (W_\mu^+ W^{-\mu} Z_\nu Z^\nu - W_\mu^+ Z^\mu W_\nu^- Z^\nu). \quad (1.3.30)$$

Finally, the abbreviated Yukawa terms in eq. (1.3.1), written in full, become

$$\begin{aligned} y_{ij} \bar{\psi}_i H \psi_j + \text{h.c.} = & - \sum_{i,j=1}^3 \left[\begin{pmatrix} \bar{\psi}_L^{\nu_i} & \bar{\psi}_L^{\ell_i} \end{pmatrix} Y_{ij}^\ell H \psi_R^{\ell_j} + \bar{\psi}_R^{\ell_j} Y_{ji}^{\ell*} H^\dagger \begin{pmatrix} \psi_L^{\nu_i} \\ \psi_L^{\ell_i} \end{pmatrix} \right] \\ & - \sum_{i,j=1}^3 \left[\begin{pmatrix} \bar{\psi}_L^{u_i} & \bar{\psi}_L^{d_i} \end{pmatrix} Y_{ij}^d H \psi_R^{d_j} + \bar{\psi}_R^{d_j} Y_{ji}^{d*} H^\dagger \begin{pmatrix} \psi_L^{u_i} \\ \psi_L^{d_i} \end{pmatrix} \right] \\ & - \sum_{i,j=1}^3 \left[\begin{pmatrix} \bar{\psi}_L^{u_i} & \bar{\psi}_L^{d_i} \end{pmatrix} Y_{ij}^u \epsilon H^* \psi_R^{u_j} + \bar{\psi}_R^{u_j} Y_{ji}^{u*} H^T \epsilon^T \begin{pmatrix} \psi_L^{u_i} \\ \psi_L^{d_i} \end{pmatrix} \right], \end{aligned} \quad (1.3.31)$$

where the indices i, j run over the three fermion generations, the $\text{SU}_L(2)$ doublets have been written explicitly, the three Yukawa coupling matrices (Y_{ij}^f) are 3×3 matrices of complex coefficients, and ϵ is the 2×2 alternating tensor (with $\epsilon_{12} = 1$),

introduced to make the up-quark-type terms gauge-invariant.¹³

After spontaneous symmetry breaking, these Yukawa terms generate mass terms for the fermions, and fermion-antifermion-Higgs interaction terms. For the quark sector, the mass terms are

$$-\frac{v}{\sqrt{2}} \sum_{i,j=1}^3 \left[\left(\bar{\psi}_L^{d_i} Y_{ij}^d \psi_R^{d_j} + \bar{\psi}_R^{d_j} Y_{ji}^{d*} \psi_L^{d_i} \right) + \left(\bar{\psi}_L^{u_i} Y_{ij}^u \psi_R^{u_j} + \bar{\psi}_R^{u_j} Y_{ji}^{u*} \psi_L^{u_i} \right) \right]; \quad (1.3.32)$$

because these are written in the basis of $SU_L(2)$ flavour-eigenstates, the masses of the physical particles are obscured. Provided that the quark masses are not degenerate, the mass matrices can be diagonalised by unitary transformations defined by four unitary matrices,

$$\bar{\Psi}_L^{d_a} = \bar{\psi}_L^{d_i} (U_L^{d\dagger})_{ia} \quad \Psi_R^{d_a} = (U_R^d)_{ai} \psi_R^{d_i} \quad (1.3.33)$$

$$\bar{\Psi}_L^{u_a} = \bar{\psi}_L^{u_i} (U_L^{u\dagger})_{ia} \quad \Psi_R^{u_a} = (U_R^u)_{ai} \psi_R^{u_i}, \quad (1.3.34)$$

such that

$$\frac{v}{\sqrt{2}} (U_L^d)_{aj} Y_{jk}^d (U_R^{d\dagger})_{kb} = M_{ab}^d = [\text{diag}(m_d, m_s, m_b)]_{ab} \quad (1.3.35)$$

$$\frac{v}{\sqrt{2}} (U_L^u)_{aj} Y_{jk}^u (U_R^{u\dagger})_{kb} = M_{ab}^u = [\text{diag}(m_u, m_c, m_t)]_{ab}. \quad (1.3.36)$$

In this basis, any other parts of the Lagrangian that are ‘diagonal’ in $SU_L(2)$ -space remain unchanged, whilst those with non-trivial $SU_L(2)$ -component are mixed. Since the generator t^3 is diagonal, this affects only the terms depending on the generators t^1 and t^2 , namely the interactions of the quarks with the physical W^+ and W^- bosons. These arise from their kinetic term,

$$i\bar{\psi}(\gamma \cdot D)\psi \ni -\frac{g}{\sqrt{2}} \begin{pmatrix} \bar{\psi}_L^{u_i} & \bar{\psi}_L^{d_i} \end{pmatrix} \begin{pmatrix} 0 & \gamma \cdot W^+ \\ \gamma \cdot W^- & 0 \end{pmatrix} \begin{pmatrix} \psi_L^{u_i} \\ \psi_L^{d_i} \end{pmatrix} \quad (1.3.37)$$

$$= -\frac{g}{\sqrt{2}} \left(\bar{\psi}_L^{u_i} (\gamma \cdot W^+) \psi_L^{d_i} + \bar{\psi}_L^{d_i} (\gamma \cdot W^-) \psi_L^{u_i} \right). \quad (1.3.38)$$

¹³In general, this could be achieved instead through a second, independent, Higgs doublet $\tilde{H}$, with opposite hypercharge to the first, $Y = -\frac{1}{2}$. Within the Standard Model, the need for a second Higgs is eliminated by relating the two through $\tilde{H} = i\sigma_2 H^*$, which has the required transformation properties and is the origin of the ϵ factors.

Transforming into the mass basis, this becomes

$$-\frac{g}{\sqrt{2}} \left(W_\mu^+ \left[\bar{\Psi}_L^{u_a} \left(U_L^u U_L^{d\dagger} \right)_{ab} (\gamma^\mu) \Psi_L^{d_b} \right] + W_\mu^- \left[\bar{\Psi}_L^{d_a} \left(U_L^d U_L^{u\dagger} \right)_{ab} (\gamma^\mu) \Psi_L^{u_b} \right] \right). \quad (1.3.39)$$

As a result, the only non-trivial dependence upon the transformation between the mass- and flavour-bases arises in the combination

$$V_{\text{CKM}} \equiv \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} := U_L^u U_L^{d\dagger}, \quad (1.3.40)$$

known as the Cabibbo–Kobayashi–Maskawa (CKM) matrix [25, 26], relating the down-type weak to the mass eigenstates as

$$\begin{pmatrix} \psi_L^d \\ \psi_L^s \\ \psi_L^b \end{pmatrix} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix} \begin{pmatrix} \Psi_L^d \\ \Psi_L^s \\ \Psi_L^b \end{pmatrix}. \quad (1.3.41)$$

With this notation, the relative strength of a weak (charged-current) interaction is given by the corresponding entry of the CKM matrix. The measurement of the entries of the CKM matrix is an active area of experimental research, as it could hint at new physics beyond the Standard Model, as will be outlined in section 1.4. The latest Particle Data Group synthesis [22] of independent direct measurements of the magnitudes of the CKM matrix entries gives

$$\begin{pmatrix} |V_{ud}| & |V_{us}| & |V_{ub}| \\ |V_{cd}| & |V_{cs}| & |V_{cb}| \\ |V_{td}| & |V_{ts}| & |V_{tb}| \end{pmatrix} = \begin{pmatrix} 0.9737 & 0.2245 & 0.0038 \\ 0.221 & 0.987 & 0.041 \\ 0.008 & 0.038 & 1.013 \end{pmatrix} \pm \begin{pmatrix} 0.00014 & 0.0008 & 0.00024 \\ 0.004 & 0.011 & 0.0014 \\ 0.0003 & 0.0011 & 0.03 \end{pmatrix}. \quad (1.3.42)$$

Other approaches extract the matrix elements from a global fit to all relevant measurements [27, 28]. It can be seen from these values that off-diagonal interactions are suppressed, penalising inter-generational quark mixing, and that most of the mixing occurs between the first and second generations. This is the basis of the Wolfenstein

parametrisation [29], which expresses the CKM matrix as a power expansion in $\lambda \approx |V_{us}|$, with parameters $A \in \mathbf{R}$ and $z = \rho + i\eta \in \mathbf{C}$,

$$V_{\text{CKM}} = \begin{pmatrix} 1 - \frac{1}{2}\lambda^2 & \lambda & A\lambda^3 z^* \\ -\lambda & 1 - \frac{1}{2}\lambda^2 & A\lambda^2 \\ A\lambda^3(1-z) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(|\lambda|^4). \quad (1.3.43)$$

A recent analysis of the experimental constraints by the CKMFitter collaboration [27] is plotted in fig. 1.3 in the Argand $(\bar{\rho}, \bar{\eta})$ -plane, where the Wolfenstein z -parameters are rescaled,

$$\bar{\rho} = \rho \left(1 - \frac{1}{2}\lambda^2\right), \quad \bar{\eta} = \eta \left(1 - \frac{1}{2}\lambda^2\right). \quad (1.3.44)$$

Within the Standard Model, the values taken by the CKM matrix dictate whether charge-parity (CP) symmetry is truly a symmetry of the electroweak Lagrangian. The charge-conjugation operation exchanges the signs of all quantum numbers, effectively exchanging particles with antiparticles, whilst the parity transformation is spatial inversion $\mathbf{x} \mapsto -\mathbf{x}$, which reverses chirality (just as ‘handedness’ is reversed in a mirror). Neither of these is a symmetry of the Standard Model; in the case of parity, the $\text{SU}_L(2)$ group explicitly distinguishes between fields of different chiralities. However, the ‘CPT theorem’ [30,31] asserts that any Lorentz-invariant and local quantum field theory with a Hermitian Hamiltonian must satisfy CPT-symmetry, in which charge-conjugation, parity inversion and time-reversal are combined. It therefore remains an open question whether the combination of charge-conjugation and parity alone is a symmetry of the Standard Model. If not, then time-reversal symmetry must also be violated.

Writing eq. (1.3.39) to explicitly project onto spinors of definite chirality, with chiral projection operators

$$P_{\text{L,R}} = \frac{1}{2} (1 \mp \gamma^5), \quad (1.3.45)$$

gives

$$-\frac{g}{\sqrt{2}} \left(W_\mu^+ \left[\bar{\Psi}^{u_a} (V_{\text{CKM}})_{ab} (\gamma^\mu P_L) \Psi^{d_b} \right] + W_\mu^- \left[\bar{\Psi}^{d_a} (V_{\text{CKM}})_{ba}^* (\gamma^\mu P_L) \Psi^{u_b} \right] \right), \quad (1.3.46)$$

which under CP-conjugation becomes

$$\begin{aligned} & -\frac{g}{\sqrt{2}} \left(W_\mu^- \left[\bar{\Psi}^{d_b} (V_{\text{CKM}})_{ab} (\gamma^\mu P_L) \Psi^{u_a} \right] + W_\mu^+ \left[\bar{\Psi}^{u_b} (V_{\text{CKM}})_{ba}^* (\gamma^\mu P_L) \Psi^{d_a} \right] \right) \\ & = -\frac{g}{\sqrt{2}} \left(W_\mu^+ \left[\bar{\Psi}^{u_a} (V_{\text{CKM}})_{ab}^* (\gamma^\mu P_L) \Psi^{d_b} \right] + W_\mu^- \left[\bar{\Psi}^{d_a} (V_{\text{CKM}})_{ba} (\gamma^\mu P_L) \Psi^{u_b} \right] \right). \end{aligned} \quad (1.3.47)$$

Comparing this with eq. (1.3.46), we see that the Standard Model violates CP-symmetry if one or more elements of V_{CKM} are not real numbers, $(V_{\text{CKM}})_{ab} \notin \mathbf{R}$.

1.4. Beyond the Standard Model?

The Standard Model is thus far remarkably effective at describing the known particles and their interactions, as shown in fig. 1.2. Any deviation from it must be subtle enough to have evaded detection at a succession of colliders. The search to verify the Standard Model at ever-increasing precisions, requiring great efforts for both theory and experiment, therefore requires motivation. Here we provide some.

Standard Model neutrinos are massless, yet the experimental observation of neutrino oscillations [33] indicates that they are massive and, as for quarks, the mass and flavour basis are not the same. The nature of the mass term to be added to the Standard Model to bring it into line with this observation is currently unknown, and will be determined by experiment to be either Majorana (if the neutrino carries no quantum numbers, in which case neutrinos will be their own antiparticle), or Dirac (like the other Standard Model fermions).¹⁴ In the latter case, it will be necessary to

¹⁴The Dirac mass term of eq. (1.2.3) has the form $m_d (\bar{\psi}_L \psi_R + \bar{\psi}_R \psi_L)$, with

$$\psi_{L,R} := P_{L,R} \psi = \frac{1}{2} (1 \mp \gamma^5) \psi, \quad (1.4.1)$$

explicitly mixing fields of different chiralities, and mass eigenstate $(\psi_L + \psi_R)$. As a result, if the neutrino field has a Dirac-type mass generated from Yukawa terms after electroweak SSB like other



Figure 1.2: ATLAS summary plot of the cross-sections measured in LHC Runs 1 and 2, and the corresponding Standard Model predictions. Taken from [32].

extend the Standard Model to include right-handed (‘sterile’) neutrinos, ν_R , which can have no gauge interactions. This is likely to be the first scheduled update to the Standard Model, but will leave other, deeper, questions unanswered.

As discussed in section 1.3, each type of fermion in the Standard Model comes in triplicate, in three generations; the quark mass eigenstates do not coincide with those of the weak interaction, leading to mixing between flavours. The CKM matrix V_{CKM} of eq. (1.3.40) must be unitary if the two bases span the same three-dimensional space. If there were another, undiscovered, generation of fermions, the restriction of the resulting 4×4 CKM matrix to the known three-generation subspace would not be unitary. This has motivated efforts to detect new physics in any deviation of the CKM matrix from unitarity. This programme has seen a remarkable advance in our determination of the CKM matrix, as shown in fig. 1.3. Recent analyses [34–36] suggest that there is a deviation of around 4σ from unitarity, called the ‘Cabibbo angle anomaly’. This could reportedly be resolved, for instance, with an additional quark or new spontaneously-broken symmetry [37], or an $\text{SU}_L(2)$ triplet fermion coupling exclusively to muons [36, 38]. This illustrates the prospect of precision electroweak measurements for both hinting at the existence of, and constraining, beyond-Standard-Model physics.

The Standard Model contains no spin-2 particle, and hence no graviton. If it is possible to describe gravity as a force mediated by a particle, like the other fundamental forces, the Standard Model is incomplete. If it is not, some deeper refinement will be required to reconcile the particle theory of the Standard Model with Einstein’s theory of general relativity. Once spacetime becomes a dynamical

fermions, a right-handed neutrino must exist. A more general mass term can be constructed,

$$\begin{pmatrix} \overline{\psi_L + \psi_L^c} & \overline{\psi_R + \psi_R^c} \end{pmatrix} \begin{pmatrix} m_L & \frac{1}{2}m_d \\ \frac{1}{2}m_d & m_R \end{pmatrix} \begin{pmatrix} \psi_L + \psi_L^c \\ \psi_R + \psi_R^c \end{pmatrix}, \quad (1.4.2)$$

where we use the shorthand $\psi_{L,R}^c := \hat{C}\gamma^0\psi_{L,R}^* = i\gamma^2\psi_{L,R}^*$ for the charge-conjugated fields and work explicitly with ‘Majorana fermion’ combinations, constructed to be invariant under $\hat{C}$ and so ‘their own antiparticle’. This contains Majorana mass terms of the form $m_R(\overline{\psi_R^c}\psi_R + \overline{\psi_R}\psi_R^c)$. Without further bosonic additions to the Standard Model $m_L = 0$ by gauge invariance of the Lagrangian, whilst m_R could be non-zero, if right-handed neutrinos exist and interact only with gravity and the Higgs mechanism.




Figure 1.3: The evolution of a global CKM fit determination of the ‘Unitarity Triangle’ between 2001 and 2019, in the rescaled Wolfenstein parametrisation of 0leq. (1.3.44). The lighter and darker shaded areas of each colour illustrate the 5% and 32% confidence limits respectively; their intersection constrains the third vertex of the triangle. The red hashed region corresponds to the 68% confidence limit. Produced by the CKMFitter Group [27].

property of the matter distribution, a theory with a very different conception of spacetime is needed, for which QFT would become a low-energy effective-theory. Candidates for such a theory famously include loop quantum gravity and string theory.

The Standard Model fails to account for the nature of dark matter, deduced to exist by cosmological observations of the rates of galactic rotations and the power-spectrum of the CMB, and dark energy, postulated by the lambda-CDM model of cosmology to explain the accelerating inflation of the universe. Nor does it explain baryon asymmetry: that our universe is apparently dominated by matter despite all SM interactions producing matter also producing antimatter in equal quantities.

As a consequence of all of these deficiencies, we know there should be physics beyond the Standard Model, but not what it might look like or where to find it. To see hints of its nature at the Large Hadron Collider, we will need to understand the experimental results we would expect the Standard Model to imply, and to be

able to calculate predictions for these measurements precisely enough to be able to identify deviations as signs of new and unexplained physical phenomena. Since the Large Hadron Collider collides hadrons, this singles out the QCD sector as requiring calculations of unprecedented precision, with tight control of all theoretical uncertainties. We therefore turn to focus upon it in greater detail.

CHAPTER 2

Quantum chromodynamics

As outlined in chapter 1, QCD is a non-abelian gauge theory based on the (unbroken) symmetry group $SU(3)_c$. Yet its realisation in nature is more complicated. The particle-field correspondence of a free quantum field theory does not, in fact, hold, and the quark and gluon fields transforming under this group do not exist as free isolated particles in nature.

In the case of QCD, they are permanently ‘confined’ in composite hadrons, and only behave as particles at short distance-scales inside high-energy collisions; the only observed isolated particles are colour singlets. This is in contrast with QED, whose corresponding fields (electrons and photons) were discovered as particles long before the theory that describes them could be fleshed out algebraically.

As a result, whilst even early perturbative predictions in QED proved extraordinarily accurate, comparable predictions in QCD require new theoretical machinery to describe the non-perturbative nature of the particle states. This takes the form of the ‘parton model’, which relates hadronic cross-sections to perturbatively-calculable partonic cross-sections through a convolution with non-perturbative ‘parton distribution functions’ (PDFs). This is described in section 2.1, whilst the calculation of the partonic cross-section is described in section 2.2.

QCD, like many other quantum field theories, contains ultraviolet divergences

occurring at high energies, which require regularisation and renormalisation, and infrared divergences occurring at low energies, which must be cancelled to permit numerical integration. These are described in section 2.3. Renormalisation leads to a residual scale-dependence of predictions on the renormalisation scale, outlined in section 2.4.

The phenomenon of confinement which leads to the parton model also leads to the impossibility of observing isolated QCD particles at colliders. Instead they combine into collimated sprays of energetic hadronic particles, called ‘jets’, introduced in section 2.5. Understanding these jets is particularly important for QCD, and for photonic final-states, as these sprays can include photons which must be distinguished from those of interest. Discussion of this will be postponed until chapter 4. Finally in section 2.6 we outline the assembly of the ideas of the rest of the chapter into a form amenable to practical calculation.

2.1. The parton model

Quantum chromodynamics describes the quark and gluon fields and their interactions, just as quantum electrodynamics describes the electron and photon. Yet unlike the latter, which are readily identified as particles, the former have never been seen in isolation. This is due to ‘confinement’.¹ Whilst the Coulomb potential of QED scales as

$$V(Q_1, Q_2, \mathbf{r}) = \frac{1}{4\pi} \frac{Q_1 Q_2}{|\mathbf{r}|} \sim \frac{1}{|\mathbf{r}|} \quad (2.1.1)$$

and so gets weaker with increasing particle separation, the corresponding QCD potential scales as $\propto \alpha_s/|\mathbf{r}|$ for small separations (less than approximately 0.1 fm), and $\propto |\mathbf{r}|$ for large separations. As a result, a quark on its way to freedom instead combines with other quarks and antiquarks produced by vacuum fluctuations to

¹ Although confinement in QCD is phenomenologically well-established, whether or not it is strictly implied by quantum Yang-Mills theory remains unclear. Proving that a quantum Yang-Mills theory implies a mass gap, which would imply confinement, is one of the Clay Mathematics Institute’s Millennium Prize problems, with a \$1 million reward.

form the free colour-singlet states we detect.

High-energy collisions of hadrons reverse this process, exposing the quarks within. Each proton at the 13 TeV LHC has a velocity $v \approx 0.9999999896 c$ and so, in the centre-of-mass frame, the incoming protons are time-dilated with a Lorentz factor $\gamma \approx 6928$. The proton's size is of the order of 1 fm, so in its rest frame the timescale for the soft gluon interactions keeping it together is of the order of 1 fm/c, and in the centre-of-mass frame, 6928 fm/c (approximately 2×10^{-20} s). If we probe the proton with a virtual particle, its lifetime is inversely proportional to its virtuality Q (by Heisenberg's uncertainty principle), and any interaction must take place within this time. As Q increases, therefore, there is a threshold above which

$$\Delta t_{\text{scatter}} \ll \Delta t_{\text{parton}}, \quad (2.1.2)$$

and the probe 'sees' a free quark, frozen within the proton and carrying some fraction of the proton momentum. This is the essence of Feynman's parton model [39], which predates its QCD context.

From the proton's perspective, there is nothing special about its configuration before the first probe from an impending collision arrives,² so we might expect this distribution to be a fundamental property of the proton. The argument is insensitive to the final-state of the scattering process, and to the nature of the probe, and so we would expect the resulting distribution of parton momenta within the proton to be universal, whatever might happen subsequently as the partons interact.

Just like the probe, the virtual gluons exchanged within the proton, holding it together, have a lifetime inversely proportional to their virtuality, and so we might expect the probability of interacting with a soft gluon to grow as the virtuality of the probe decreases and the timescale of the scattering grows. Accordingly, we should expect some dependence of the distribution of partonic momenta on the scale of the

² In general, an interaction will not be with the first probe. One might therefore worry that, for instance, soft gluons emitted from one hadron might change the distribution of partons in the other before any hard probe arrives. It can be shown [40] by analogy with electrodynamics that such contributions contribute, in the worst case, to 'higher twist' terms in eq. (2.1.3), suppressed by a factor of $1/s^2$, and so do not spoil the general argument.

probe.

From first principles, therefore, the distribution of partons exposed within a hadron H might be conjectured to follow some set of probability distribution functions $\{f_i^H(x_i; Q^2)\}_{\{i \in H\}}$, representing the probability that parton i is found carrying (longitudinal) fraction x_i of the parent momentum P_H when probed at virtuality Q^2 . But we should not expect to be able to calculate these functions perturbatively, since we know that at low scales perturbation theory will break down.

If there can be no interference between the interactions of partons amongst themselves, and the interactions of a parton with the probe, we can neglect the interference terms that would arise from the quantum-mechanical calculation with probability amplitudes, and revert to a classical calculation with probabilities [41].³ This leads to a streamlined way of calculating hadronic cross-sections perturbatively without a perturbative description of the hadron. We can compute partonic cross-sections for each possible configuration of partons inside the colliding hadrons, modify them by the probability of finding those partons with that momentum configuration, and sum over all possible configurations (by integrating over momentum fractions, and summing over partons).

Although this can be motivated on physical grounds as above, by comparing timescales, formal proofs are difficult to obtain. So far they are limited to inclusive final-states produced via the Deep Inelastic Scattering (DIS) of a lepton from a hadron [41, 43, 44], and via the Drell-Yan process, in which a quark from one hadron and an antiquark from another annihilate to produce a Z -boson, which decays to leptons [41, 44, 45]. In other cases we rely on a factorisation ansatz.

³ This non-trivial assumption is known as the quantum mechanical ‘incoherence’ of the parton model. The interference terms we neglect as a result include cross-terms between diagrams in which the scattering of, say, a quark with momentum fraction x_i interferes with the scattering of a quark with momentum fraction x'_i , both from the same hadron [42]. As a result, we think of the high-energy scattering as occurring between a single parton from each hadron, each with definite partonic identity and definite momentum. Proving that this is legitimate amounts to a proof of factorisation.

2.1.1. Factorisation

This intuition is embodied in the collinear factorisation⁴ equation for a hadronic cross-section:

$$\begin{aligned} \sigma_{AB \rightarrow X} = & \sum_{a,b} \int_0^1 dx_a \int_0^1 dx_b f_a^A(x_a; \mu_F^2) f_b^B(x_b; \mu_F^2) \int d\Phi_{ab \rightarrow \hat{X}} \frac{d\hat{\sigma}_{ab \rightarrow \hat{X}}(\Phi_{ab \rightarrow \hat{X}}; \mu_F^2)}{d\Phi_{ab \rightarrow \hat{X}}} \\ & + \mathcal{O}\left(\frac{\Lambda_{\text{QCD}}^2}{s}\right), \end{aligned} \quad (2.1.3)$$

where μ_F is the ‘factorisation scale’ and represents the scale at which the corresponding hadron is probed. This equation is illustrated diagrammatically in fig. 2.1. This is not exact, as indicated by the remainder term, but is instead the leading term in some power-series expansion, with power correction terms (conventionally called ‘higher-twist’) suppressed by powers of ratios of scales which are small for high-energy collisions. Here, the identified final state of the hadronic process X is implicitly assumed to originate as the partonic final state $\hat{X}$ in the hard scattering, but in general, this may not be the case: an identified final-state jet of mesons and baryons might arise from a partonic final-state of quarks or gluons. In this case a transition function $F(\hat{X} \rightarrow X; \mu_F^2, \mu_f^2)$ can be included, accounting for the probability that the high-energy and short-range partonic final-state $\hat{X}$ gives the long-range hadronic observable X after non-perturbative evolution to some lower scale μ_f .

The parton distribution functions f_i^H cannot in fact be probability distribution functions, which would require them to integrate to 1; instead they are number

⁴ In collinear factorisation we make the assumption that a parton carries a longitudinal fraction of the parent hadron’s momentum, with zero transverse momentum. It is also possible to define factorisation more generally, with parton distribution functions which additionally depend on the transverse momentum of the parton. The PDFs for collinear factorisation then correspond to the integral of these ‘transverse-momentum dependent’ (TMD) PDFs over the transverse momentum $\mathbf{k}_T$. TMD factorisation applies for $q_T \ll Q$, and its accuracy decreases as q_T increases, whilst collinear factorisation degrades as $q_T \rightarrow 0$ (and, uncorrected, leads to unphysical singularities in the limit). We focus exclusively on collinear factorisation here; see [44] for more detail.




Figure 2.1: Schematic representation of the hadronic scattering process $A + B \rightarrow X$, illustrating the ‘factorisation’ relation of the hadronic cross-section into the convolution of the parton density functions $f_a^A(x_a)$ and $f_b^B(x_b)$ with the partonic cross-section for the hard scattering, $\hat{\sigma}_{ab}$. The white box to the right represents the subsequent evolution of the partonic final-state to lower scales, in the parton shower and through hadronisation.

densities, and must be normalised to obey conservation of momentum,

$$\sum_i \int_0^1 x f_i^H(x; \mu_F^2) dx = 1, \quad (2.1.4)$$

and to reproduce the known valence quark properties of the hadron, e.g.:

$$\begin{aligned} \int_0^1 [f_u^p(x; \mu_F^2) - f_{\bar{u}}^p(x; \mu_F^2)] dx &= 2, \\ \int_0^1 [f_d^p(x; \mu_F^2) - f_{\bar{d}}^p(x; \mu_F^2)] dx &= 1, \\ \int_0^1 [f_s^p(x; \mu_F^2) - f_{\bar{s}}^p(x; \mu_F^2)] dx &= 0. \end{aligned} \quad (2.1.5)$$

2.1.2. DGLAP evolution

We can apply perturbation theory to our understanding of the parton distribution functions $f_i^H(x; \mu_F^2)$ to derive further properties they must satisfy. There can be no residual dependence of the left-hand-side of eq. (2.1.3), which corresponds to a prediction for a measured quantity, on the theoretical parameter μ_F . The change in

a PDF as μ_F changes must therefore be compensated by a balancing change in the partonic cross-section, to leave the convolution μ_F -independent.

This balance leads to the Dokshitzer–Gribov–Lipatov–Altarelli–Parisi (DGLAP) equations, which also arise independently from the renormalisation of the operators contributing to the formally-defined parton distribution functions within the Operator Product Expansion formalism (for details, see [9]). The DGLAP equations are a set of $2n_f + 1$ coupled integro-differential equations relating the evolution of the PDFs for each parton to the space-like⁵ QCD splitting functions:

$$\begin{aligned}\frac{\partial f_{q_i}(x, \mu_F^2)}{\partial \ln \mu_F^2} &= \frac{\alpha_s(\mu_F^2)}{2\pi} \int_x^1 \frac{dz}{z} \left[f_{q_i}\left(\frac{x}{z}, \mu_F^2\right) P_q^q(z) + f_g\left(\frac{x}{z}, \mu_F^2\right) P_q^g(z) \right] \\ \frac{\partial f_{\bar{q}_i}(x, \mu_F^2)}{\partial \ln \mu_F^2} &= \frac{\alpha_s(\mu_F^2)}{2\pi} \int_x^1 \frac{dz}{z} \left[f_{\bar{q}_i}\left(\frac{x}{z}, \mu_F^2\right) P_q^q(z) + f_g\left(\frac{x}{z}, \mu_F^2\right) P_q^g(z) \right] \\ \frac{\partial f_g(x, \mu_F^2)}{\partial \ln \mu_F^2} &= \frac{\alpha_s(\mu_F^2)}{2\pi} \int_x^1 \frac{dz}{z} \left[\sum_{i=1}^{n_f} \left[f_{q_i}\left(\frac{x}{z}, \mu_F^2\right) + f_{\bar{q}_i}\left(\frac{x}{z}, \mu_F^2\right) \right] P_g^q(z) \right. \\ &\quad \left. + f_g\left(\frac{x}{z}, \mu_F^2\right) P_g^g(z) \right]\end{aligned}\quad (2.1.6)$$

This has a natural interpretation: a parton b resolved at scale μ_F could have been produced from the collinear splitting of a parton a resolved at scale $\mu_F + \delta\mu_F$, and so we have a Markov-like process computable as the sum over all possible identities and longitudinal momentum fractions of the intermediate parton, with the splitting functions $P_b^a(z)$ taking the role of the Markov kernel and encoding the transition probability from parton a to collinear parton b . This idea is illustrated diagrammatically in fig. 2.2.

The splitting functions are calculable within perturbative QCD as the ratio of the matrix elements $|\mathcal{M}_{a \rightarrow b \rightarrow X}|^2$ and $|\mathcal{M}_{b \rightarrow X}|^2$, in the collinear limit where $a \parallel b$, up to phase-space factors. The method is general, leading to universal, process-independent functions computable order-by-order by the usual methods of perturbative QCD. Remarkably, they have recently been calculated to three-loops [48, 49]. To leading

⁵ Beyond leading-order, the ‘space-like’ and ‘time-like’ splitting functions for initial- and final-state collinear factorisation (named according to the four-momentum q transferred by the virtual probe) differ, although they remain related through analytic continuation up to NNLO [46, 47]. This subtlety will be revisited in the context of final-state factorisation, fragmentation, in section 4.1.




Figure 2.2: Diagrammatic representation of the DGLAP equations for the scale-evolution of the parton distribution functions. The antiquark-to-gluon contributions have been included in the sum over quark diagrams for brevity. Here $t = \mu_F^2$, so $t \partial F / \partial t \equiv \partial F / \partial (\ln \mu_F^2)$.

order in α_s ,⁶ the splitting functions $P_b^a(z)$ are [50]:

$$P_q^{q(0)}(z) = C_F \left[\frac{1+z^2}{(1-z)_+} + \frac{3}{2} \delta(1-z) \right] \quad (2.1.7)$$

$$P_g^{q(0)}(z) = C_F \left[\frac{1+(1-z)^2}{z} \right] \quad (2.1.8)$$

$$P_q^{g(0)}(z) = T_F [z^2 + (1-z)^2] \quad (2.1.9)$$

$$P_g^{g(0)}(z) = 2C_A \left[\frac{z}{[1-z]_+} + \frac{1-z}{z} + z(1-z) \right] + \frac{1}{6} (11C_A - 4n_f T_F) \delta(1-z) \quad (2.1.10)$$

where divergences are regularised using the ‘+’-prescription, which defines a distribution by its action on a test function f such that

$$\int_0^1 f(z) [g(z)]_+ dz := \int_0^1 [f(z) - f(1)] g(z) dz. \quad (2.1.11)$$

This cancels the divergences of $g(z)$ at $z = 1$, provided that $f(z)$ is sufficiently smooth there.

⁶ The coupling-constant renormalisation scale in the perturbative expansion of the splitting functions is set to the factorisation scale, so the expansion is in $\alpha_s(\mu_F^2)$. This will be discussed further in section 2.3.

There is an important caveat to eq. (2.1.3). The PDFs and the partonic cross-section, naïvely defined, are each independently divergent. This will necessitate the discussion of regularisation and renormalisation in sections 2.3.1.1 and 2.3.1.2. Once an appropriate renormalisation procedure is defined, eq. (2.1.3) holds exactly as written for the renormalised PDFs and partonic cross-sections, whilst the DGLAP equations eq. (2.1.6) can then be identified as the renormalisation group equations in the unphysical mass-factorisation scale. This will be revisited in section 2.3.

2.1.3. Testing the parton model

Due to confinement it is impossible to test perturbative QCD directly. Our success at using perturbative QCD to make predictions for collider experiments is therefore as much a test of the parton model and factorisation as it is of perturbative QCD itself.

The central testable claim of the parton model is that there is a universal set of parton distribution functions for each hadron for which eq. (2.1.3) holds, for all collisions involving that hadron, all centre-of-mass energies, all probe particles and all identified final states.

Although it may be possible to calculate the parton distribution functions from first principles using lattice QCD, this is currently beyond our capabilities. Current approaches focus on computing the lowest three Mellin moments,⁷ which constrain but do not fully specify the momentum-fraction dependence of the PDFs [51].

Instead, PDFs are determined from a global best-fit of measurements to theory predictions, in which a PDF is parametrised at some low scale and the parametrisation evolved using the DGLAP equations to the relevant factorisation scale. The

⁷ The Mellin transformation of a function $f(x)$ is defined as the integral transform with kernel $K(s, x) = x^{s-1}$, i.e.

$$\mathcal{M}[f](s) = \int_0^1 x^{s-1} f(x) dx,$$

with the n -th Mellin moment defined as $\mathcal{M}[f](n)$ for $n \in \mathbf{N}$. This is especially useful for studying solutions to the DGLAP equations, since it maps the relevant convolutions of functions onto products of their Mellin transforms.

resulting comparison of theory to experiment constrains the parameters, which are then determined by the minimisation of some statistical goodness-of-fit parameter, such as χ^2 for (weighted) least-squares minimisation. The existence of a PDF that simultaneously satisfies all the constraints, to a degree commensurate with the necessary experimental and theoretical errors, would confirm the applicability of the parton model and our understanding of perturbative QCD.

However, the combination of experimental errors, theory uncertainties from missing higher-orders in the hard-scattering calculations, theory uncertainties from the fitting and optimisation procedure, and extrapolation uncertainties in kinematical regions poorly constrained by experimental data, are all difficult to calculate and assess. This can lead to a considerable discrepancy even between PDF sets using the same input data but different methodologies. Different groups tackle these challenges differently, with some, such as HERAPDF [52] choosing to use a constrained but consistent set of data from a single experiment (in this case, the $e^\pm p$ collider HERA), giving explicit uncertainties where the PDFs are unconstrained, and others seeking to maximise the kinematic range of the input data by using inputs from fixed-target experiments, HERA, the Tevatron $p\bar{p}$ collider, and the LHC (such as NNPDF [53], MMHT [54] and CTEQ [55]). Uncertainties then lead to a deterioration in the quality-of-fit parameter. As a result, it is hard to rule out the possibility that new, beyond-Standard-Model physics inadequately described by QCD could simply hide in the uncertainties in our knowledge of the PDFs.

Improving this situation is one of the main motivations for the Large Hadron-electron Collider (LHeC) [56], a programme to upgrade the LHC concurrently with the planned High-Luminosity LHC upgrade, to allow its simultaneous operation as a $e^\pm p$ and a $p\bar{p}$ collider. Such a collider would enable the precise determination of PDFs through charged- and neutral-current DIS alone, across a wide range of longitudinal momentum fractions x (from approximately 10^{-6} to 0.9) and energy scales Q^2 (up to $1.7 \times 10^6 \text{ GeV}^2$). Predictions using these PDFs would then provide a precision test of factorisation and the parton model for proton-proton collisions,




Figure 2.3: Proton PDFs at two different scales, emphasising the effect of DGLAP evolution. Note the dominance of the gluon PDF for small x , and that of the valence up and down quark distributions for large x . This set of PDFs was produced by the CTEQ collaboration in 2019 [57] from fits to HERA and LHC data, for convolution with NNLO partonic calculations. The uncertainty bands are not shown.

by ensuring that signs of factorisation-violation could not be misattributed to an inadequate knowledge of the proton's structure.

2.2. Partonic cross-sections

We now turn to the computation of the partonic cross-section arising within the factorisation equation eq. (2.1.3), $\hat{\sigma}_{ab \rightarrow \hat{X}}$. The perturbative expansion in the QCD coupling constant gives

$$\begin{aligned} d\hat{\sigma}_{ab \rightarrow X} = & \left(\frac{\alpha_s}{2\pi}\right)^m d\hat{\sigma}_{ab}^{\text{LO}} + \left(\frac{\alpha_s}{2\pi}\right)^{m+1} d\hat{\sigma}_{ab}^{\text{NLO}} + \left(\frac{\alpha_s}{2\pi}\right)^{m+2} d\hat{\sigma}_{ab}^{\text{NNLO}} \\ & + \mathcal{O}(\alpha_s^{m+3}), \end{aligned} \quad (2.2.1)$$

where we can identify terms of this expansion with that of $|M_{i \rightarrow f}|^2$, so that

$$\begin{aligned} |M|^2 = & \left(\frac{\alpha_s}{2\pi}\right)^m |M|_{\text{LO}}^2 + \left(\frac{\alpha_s}{2\pi}\right)^{m+1} |M|_{\text{NLO}}^2 + \left(\frac{\alpha_s}{2\pi}\right)^{m+2} |M|_{\text{NNLO}}^2 \\ & + \mathcal{O}(\alpha_s^{m+3}). \end{aligned} \quad (2.2.2)$$

Here ‘LO’ stands for ‘leading-order’, indicating the first non-trivial contribution to the perturbative expansion, and the coefficients of higher powers are called ‘next-to-leading-order’ (NLO), ‘next-to-next-to-leading-order’ (NNLO), and subsequently $N^k\text{LO}$ for brevity. Each additional factor of α_s corresponds to two additional vertices, each with vertex factor $g_s = \sqrt{4\pi\alpha_s}$, in the Feynman diagram expansion.

In fact we consider inclusive final states $f + X$, so must sum over the matrix elements corresponding to additional emissions, which we will regard as additional contributions to $d\hat{\sigma}^{N^k\text{LO}}$ at each order, decomposed according to whether additional vertices attach to additional external particles, or additional internal propagators to create loops, which we call ‘real’ and ‘virtual’ contributions respectively. In the former case, we can write the n -particle amplitude as

$$M_n = g_s^m M_n^{(0)} + g_s^{m+2} M_n^{(1)} + g_s^{m+4} M_n^{(2)} + \mathcal{O}(g_s^{m+6}), \quad (2.2.3)$$

as each additional internal propagator must connect at both ends,⁸ whilst for the latter case we can write the corresponding expansion for the $(n+1)$ - and $(n+2)$ -particle external state, at one and two orders higher in g_s ,

$$M_{n+1} = g_s^{m+1} M_{n+1}^{(0)} + g_s^{m+3} M_{n+1}^{(1)} + \mathcal{O}(g_s^{m+5}), \quad (2.2.4)$$

$$M_{n+2} = g_s^{m+2} M_{n+2}^{(0)} + g_s^{m+4} M_{n+2}^{(1)} + \mathcal{O}(g_s^{m+6}). \quad (2.2.5)$$

As a result,

$$\begin{aligned} |M_n|^2 &= g_s^{2m} |M_n^{(0)}|^2 \\ &\quad + g_s^{2m+2} [M_n^{(1)\dagger} M_n^{(0)} + M_n^{(0)\dagger} M_n^{(1)}] \\ &\quad + g_s^{2m+4} \left[|M_n^{(1)}|^2 + (M_n^{(2)\dagger} M_n^{(0)} + M_n^{(0)\dagger} M_n^{(2)}) \right] \\ &\quad + \mathcal{O}(g_s^{2m+6}), \end{aligned} \quad (2.2.6)$$

where each line is called the ‘Born’ (B), ‘virtual’ (V) and ‘double virtual’ (VV)

⁸ Creating a new vertex or adding a particle to a vertex where possible both carry a factor of g_s in the Feynman rules for QCD.

contribution respectively;

$$\begin{aligned}
|M_{n+1}|^2 &= g_s^{2m+2} |M_{n+1}^{(0)}|^2 \\
&+ g_s^{2m+4} [M_{n+1}^{(1)\dagger} M_{n+1}^{(0)} + M_{n+1}^{(0)\dagger} M_{n+1}^{(1)}] \\
&+ \mathcal{O}(g_s^{2m+6}),
\end{aligned} \tag{2.2.7}$$

where each line is called the ‘real’ (R) and ‘real-virtual’ (RV) contribution respectively; and

$$|M_{n+2}|^2 = g_s^{2m+4} |M_{n+2}^{(0)}|^2 + \mathcal{O}(g_s^{2m+6}), \tag{2.2.8}$$

called the ‘double-real’ (RR) contribution. From this decomposition, we can see that every contribution has an even power of g_s , and can equate coefficients at each power in $\alpha_s \propto g_s^2$ to read off the Feynman diagrams contributing to each of $d\hat{\sigma}^{\text{LO}}$, $d\hat{\sigma}^{\text{NLO}}$ and $d\hat{\sigma}^{\text{NNLO}}$.

There is a final subtlety. For Standard Model calculations, when f includes non-QCD particles (for instance, photons), the m for which the first term in the perturbative series is non-trivial differs between partonic channels. This leads to the leading-order of one partonic channel entering the perturbative series at the same order in α_s as a higher order of another, as defined above. In such cases we define LO, NLO, NNLO, etc. according to the lowest m for any partonic channel.

2.3. Divergences in QCD

In QCD, as in other quantum field theories, certain diagrams in the Feynman expansion are divergent due to loop integrals such as the ‘bubble’ integral,⁹

$$\int_0^\Lambda \frac{d^4k}{(2\pi)^4} \frac{1}{[k^2 + i0][(k+p)^2 + i0]} \sim \log \Lambda \xrightarrow{\Lambda \rightarrow \infty} \infty, \quad (2.3.1)$$

where we have introduced a cutoff Λ to capture the divergence as $k^2 \rightarrow \infty$. Because this divergence arises from the integral over the large-momentum modes, it is called an ‘ultraviolet’ (UV) divergence. These are removed by ‘regularisation’, in which an unphysical scale is introduced to capture the divergent limit as a divergent function of this new parameter. This is summarised in section 2.3.1.1. This regularisation procedure removes any predictive power from the theory and so necessitates its restoration through the re-normalisation of the fields and couplings, outlined in section 2.3.1.2.

The bubble integral in eq. (2.3.1) also contains a divergence associated with low-momentum modes, called an ‘infrared’ divergence. These arise in massless quantum field theories, including massless QCD, from the singular behaviour of an internal massless propagator when the four-momentum goes ‘on-shell’. Closely-related divergences arise in the real-emission diagrams in the limits associated with low-energy and collinear particles. Crucially, these divergences cancel in the sum over all contributions at a given order of the perturbative series, a fact that will be exploited in section 2.6 to simplify numerical calculations.

⁹ The $+i0$ term here represents the inclusion of a small imaginary part, $+i\delta$, and the subsequent $\delta \rightarrow 0$ limit. This can be thought of as an infinitesimal complex-plane rotation, specifying the location of the poles in the denominator relative to the k^0 -axis. In general, this term enforces the desired causality of the propagator by ensuring that the k^0 integral along the real axis is a deformation of the ‘Feynman contour’ and so fixes the boundary conditions of the Green’s function accordingly. For loop integrals like this one, it also regularises the divergence that would otherwise occur at $k = -p$, allowing the unambiguous Wick-rotation of the k_0 -coordinate into Euclidean space, $k_E^0 := e^{-i\theta} k^0$ for $\theta = \pi/2$, without passing through any poles. The integral can then be rewritten using the Euclidean metric, expressed in generalised spherical coordinates, and evaluated. In subsequent expressions the $+i\delta$ term will be omitted, but should be understood to be present. It can be restored with the consistent replacement $k^2 \rightarrow k^2 + i\delta$.

2.3.1. Ultraviolet divergences

To properly define the integral in this limit, the divergence must be regularised (e.g. as above by introducing a cutoff), capturing the infinite contribution. Any two schemes differing by a constant would work equally well for this, so a choice must be made that restores connection with physical quantities. These two steps are called regularisation and renormalisation, and introduce a new dependence on an unphysical scale μ_R .

2.3.1.1. Regularisation

The prevailing technique for regularising divergent Feynman integrals is ‘dimensional regularisation’, in which the integral is formulated in d -dimensions, evaluated where it converges (typically using generalised expressions for the volume of a d -dimensional hypersphere), and analytically-continued to values of d for which the integral diverges and so is otherwise undefined.

In dimensional regularisation, defining $\varepsilon = \frac{1}{2}(4 - d)$, the previous example eq. (2.3.1) becomes, through Feynman parametrisation,

$$\int \frac{d^d k}{(2\pi)^d} \frac{1}{k^2(k+p)^2} = \frac{\Gamma(\varepsilon)}{(4\pi)^{2-\varepsilon}} (-p^2)^{-\varepsilon} B(1-\varepsilon, 1-\varepsilon), \quad (2.3.2)$$

where $B(x, y)$ is the Beta function.¹⁰ Here the divergence of the integral for $d = 4$ (equivalently, $\varepsilon = 0$) has been captured by the $\Gamma(\varepsilon)$ factor, where $\Gamma(z)$ is the Gamma function, which has Laurent expansion

$$\Gamma(z) = \frac{1}{z} - \gamma_E + \frac{1}{2} \left(\gamma_E^2 + \frac{\pi^2}{6} \right) z + \mathcal{O}(z^2), \quad (2.3.5)$$

¹⁰The Beta function is defined by

$$B(x, y) := \int_0^1 t^{x-1} (1-t)^{y-1} dt \quad \text{Re } x, \text{Re } y > 0 \quad (2.3.3)$$

$$\equiv \frac{\Gamma(x)\Gamma(y)}{\Gamma(x+y)}. \quad (2.3.4)$$

This Gamma-function identity provides the analytic continuation of the Beta function to the remainder of the complex plane, where the integral definition eq. (2.3.3) fails to converge.

where $\gamma_E \approx 0.577$ is the Euler–Mascheroni constant, and so we can verify that our evaluated integral indeed diverges as $\varepsilon \rightarrow 0$.

Since the action S must be dimensionless, adjusting the dimension of integration from d^4x to $d^d x$ must change the mass-dimension of the Lagrangian density $\mathcal{L}$ from 4 to d . The kinetic terms in the Lagrangian for a given field determine their dimension, and so we can deduce the dimension of each coupling from the interaction terms. For QCD the three-gluon vertex implies that we require

$$g \rightarrow g\mu^\varepsilon, \quad (2.3.6)$$

which ensures that the coupling-constant remains dimensionless, but introduces the expected dependence of the result on the scale μ .

2.3.1.2. Renormalisation

In QCD, only finitely many types of Feynman diagram are divergent.¹¹ As a result, we need only finitely many parameters to absorb these infinities and assign physical values afterwards. We exploit the freedom to choose the normalisation of the fields, and re-normalise them multiplicatively. We then call the parameters and fields appearing directly in the Lagrangian ‘bare’, and define the new renormalised parameters to be related by multiplicative factors, such as

$$\psi_0(x) = \sqrt{Z_2} \psi(x), \quad G_0^\mu(x) = \sqrt{Z_3} G^\mu(x), \quad g_0 = Z_g \mu^\varepsilon g. \quad (2.3.7)$$

The resulting Lagrangian density can then be considered as

$$\mathcal{L} = \mathcal{L}_{\text{bare}} + \mathcal{L}_{\text{c.t.}} \quad (2.3.8)$$

where the counterterm Lagrangian contains the deviation of each Z_i from the bare value, $\delta_i = Z_i - 1$ (the interaction counterterms are slightly more complicated, as they correspond to products of fields). Renormalisation then corresponds to choosing

¹¹The fundamental divergent 1PI diagrams can be counted and classified combinatorially, but can appear as subdiagrams of others, making them divergent as a result.

a non-zero value for each δ_i .

The constraints imposed to fix each counterterm δ_i are called a ‘renormalisation scheme’. The conventional scheme for QCD is modified-minimal-subtraction ($\overline{\text{MS}}$). Minimal subtraction (MS) is defined by the constraint that the counterterms contain only the principal part of the Laurent series arising from the Feynman diagram calculation, subtracting the ε -poles but no more. Modified minimal-subtraction is defined by minimal subtraction with an additional inverse factor included for each loop,

$$\begin{aligned} C_\varepsilon &= (4\pi)^\varepsilon e^{-\varepsilon\gamma_E} \\ &= 1 + \varepsilon (\log 4\pi - \gamma_E) + \mathcal{O}(\varepsilon^2) \\ &= \frac{(4\pi)^\varepsilon}{\Gamma(1-\varepsilon)} + \mathcal{O}(\varepsilon^2), \end{aligned} \tag{2.3.9}$$

to remove this universal geometric factor, which arises from the angular integration in d -dimensions that is performed for each loop momentum.¹² The $\overline{\text{MS}}$ renormalisation factors have recently been calculated in full to the five-loop level [58].

As an example, the NLO loop corrections to the quark propagator can be computed to fix (with gauge parameter ξ) [10]

$$Z_2^{\overline{\text{MS}}} = 1 - \frac{\alpha_s}{4\pi} \frac{C_\varepsilon}{\varepsilon} \xi C_F + \mathcal{O}(\alpha_s^2) \tag{2.3.10}$$

while those for the gluon propagator give

$$Z_3^{\overline{\text{MS}}} = 1 - \frac{\alpha_s}{4\pi} \frac{C_\varepsilon}{\varepsilon} \left(\frac{4}{3} T_F n_f + \left(\frac{\xi}{2} - \frac{13}{6} \right) C_A \right) + \mathcal{O}(\alpha_s^2) \tag{2.3.11}$$

and the loop corrections to the quark-gluon vertex can be evaluated and imply that

$$Z_{q\bar{q}g}^{\overline{\text{MS}}} = 1 - \frac{\alpha_s}{4\pi} \frac{C_\varepsilon}{\varepsilon} \left(\frac{3}{4} C_A + \xi \left(\frac{1}{4} C_A + C_F \right) \right) + \mathcal{O}(\alpha_s^2). \tag{2.3.12}$$

¹²This factor can be identified in the bubble integral in eq. (2.3.2) using the Beta-function identity of eq. (2.3.4), which implies

$$B(1-\varepsilon, 1-\varepsilon) = \frac{\Gamma(1-\varepsilon)^2}{\Gamma(2-2\varepsilon)}.$$

Taken together, these three results imply

$$\begin{aligned} Z_g^{\overline{\text{MS}}} &= \frac{Z_{q\bar{q}g}^{\overline{\text{MS}}}}{Z_2^{\overline{\text{MS}}} \sqrt{Z_3^{\overline{\text{MS}}}}} \\ &= 1 - \frac{\alpha_s}{4\pi} \frac{C_\varepsilon}{\varepsilon} \left(\frac{11}{6} C_A - \frac{2}{3} T_F n_f \right) + \mathcal{O}(\alpha_s^2). \end{aligned} \quad (2.3.13)$$

Although it would be possible for the $\{Z_i\}$ to be overconstrained, the pole pieces computed from the various loop corrections are all consistent with each other, a remarkable and important check on the validity of the renormalisation procedure.

2.3.1.3. The running coupling

The bare coupling constant g_0 cannot depend on the scale introduced to make it dimensionless under dimensional regularisation. This is a remarkably powerful statement. Using the expansion parameter $\alpha_s = g_s^2/4\pi$, this implies

$$\begin{aligned} \beta(\alpha_s) &:= \frac{d\alpha_s(\mu)}{d \ln \mu^2} \\ &= -\varepsilon \alpha_s - 2\alpha_s \frac{1}{Z_g} \frac{\partial Z_g}{\partial \alpha_s} \frac{d\alpha_s(\mu)}{d \ln \mu^2} \\ &\xrightarrow{\varepsilon \rightarrow 0} -2\alpha_s \frac{1}{Z_g} \frac{\partial Z_g}{\partial \alpha_s} \beta(\alpha_s), \end{aligned} \quad (2.3.14)$$

and so in the $\overline{\text{MS}}$ scheme, the scale evolution of the coupling is determined by the counterterm $Z_g^{\overline{\text{MS}}}$ for the renormalisation of the coupling constant, given in eq. (2.3.13). As a result, in this renormalisation scheme the scale-dependence of the strong coupling must satisfy the differential equation

$$\frac{d\alpha_s(\mu)}{d \ln \mu^2} = -\frac{\alpha_s^2}{4\pi} \left(\frac{11}{3} C_A - \frac{4}{3} T_F n_f \right) + \mathcal{O}(\alpha_s^2), \quad (2.3.15)$$

or, as $df(x)/d(\ln x^2) = \frac{x}{2} df/dx$,

$$\frac{d\alpha_s(\mu)}{d\mu} = -\mu^{-1} \frac{\alpha_s^2}{2\pi} \left(\frac{11}{3} C_A - \frac{4}{3} T_F n_f \right) + \mathcal{O}(\alpha_s^2). \quad (2.3.16)$$

Higher-order terms in the expansion of the beta function,

$$\beta(\alpha_s) = -\alpha_s \sum_{n=0}^{\infty} \beta_n \left(\frac{\alpha_s}{4\pi} \right)^{n+1}, \quad (2.3.17)$$

have been calculated up to $n = 4$ [59, 60], requiring the evaluation of five-loop diagrams.

If $\alpha_s(\mu)$ is small, the higher-order terms can be discarded and eq. (2.3.17) can be integrated directly, to give

$$\alpha_s(\mu) = \frac{\alpha_s(Q)}{1 + \beta_0 \frac{\alpha_s(Q)}{4\pi} \log\left(\frac{\mu^2}{Q^2}\right)}, \quad (2.3.18)$$

which vanishes as $\mu \rightarrow \infty$. This is called ‘asymptotic freedom’, and justifies our application of perturbative QCD to ‘hard’, high-energy scattering events with large characteristic momentum scales.

In QCD, $C_A = 3$ and by convention $T_F = \frac{1}{2}$, and so for $n_f < \frac{33}{2}$ the sign of the leading term of the beta-function is negative. This is the case in nature, where there are only six identified flavours of quarks.¹³ This is in contrast to QED, whose beta function has a leading term that is always positive,

$$\frac{d\alpha_{\text{em}}(\mu)}{d \ln \mu^2} = \frac{\alpha_{\text{em}}^2}{4\pi} \left(\frac{4}{3} n_f \right) + \mathcal{O}(\alpha_{\text{em}}^2). \quad (2.3.19)$$

As a result, the strong coupling constant $\alpha_s(\mu)$ must diverge as $\mu \rightarrow 0$, while the electromagnetic coupling constant of QED tends to a constant, the ‘fine-structure’ constant. This suggests that perturbative QCD breaks down in this limit, consistent with our expectations of confinement from the parton model.

The expected scale-dependence of the coupling has been confirmed experimentally, as shown in fig. 2.4.

2.3.1.4. Renormalisation of matrix elements and cross-sections

The Feynman rules are expressed in terms of the bare coupling g_0 , but the renormalisation of the coupling mixes powers of g_0 through Z_g , which is itself a perturbative expansion in g . The decomposition of the cross-section order-by-order in α_s therefore changes in the transition between the bare and the renormalised couplings.

¹³ In practice quarks which are heavy relative to the scale are neglected, as their contribution is suppressed by powers of m^{-1} for $Q \ll m_q$.




Figure 2.4: PDG summary of experimental measurements of the strong coupling α_s as a function of the energy scale Q , as of 2020. Taken from [22].

This can be derived for a general perturbative expansion as

$$\sum_{k=0}^{\infty} a_k \left(\frac{\alpha_s(\mu)}{2\pi} \right)^{m+k} = \sum_{k=0}^{\infty} a_k \left(\frac{\alpha_s^0}{2\pi} \mu^{-2\varepsilon} \right)^{m+k} \left(Z_g^{-2} \right)^{k+m} \quad (2.3.20)$$

$$= \left(\frac{\alpha_s^0}{2\pi} \mu^{2\varepsilon} \right)^m \left[a_0 + \left(\frac{\alpha_s^0}{2\pi} \mu^{2\varepsilon} \right) \left[a_1 - 2a_0 m Z_g^{(1)} \right] \right. \quad (2.3.21)$$

$$+ \left(\frac{\alpha_s^0}{2\pi} \mu^{2\varepsilon} \right)^2 \left[a_2 - 2a_1(m+1)Z_g^{(1)} + 2a_0 m \left(\left(m + \frac{1}{2} \right) Z_g^{(1)2} - Z_g^{(2)} \right) \right] \left. \right]$$

$$+ \mathcal{O}(\alpha_s(\mu)^3)$$

where $Z_g^{(k)}$ denotes the coefficient of $(\alpha_s/2\pi)^k$ in the expansion of Z_g .

This can be applied directly to the perturbative expansion of the cross-section, giving for eq. (2.2.1) with $m = 0$, the transitions

$$d\hat{\sigma}^{\text{LO}} = d\hat{\sigma}^{0,\text{LO}} \quad (2.3.22)$$

$$d\hat{\sigma}^{\text{NLO}} = d\hat{\sigma}^{0,\text{NLO}} \quad (2.3.23)$$

$$d\hat{\sigma}^{\text{NNLO}} = d\hat{\sigma}^{0,\text{NNLO}} - 2Z_g^{(1)} d\hat{\sigma}^{0,\text{NLO}}$$

$$= d\hat{\sigma}^{0,\text{NNLO}} - \frac{C_\varepsilon}{\varepsilon} \beta_0 d\hat{\sigma}^{0,\text{NLO}}, \quad (2.3.24)$$

where $d\hat{\sigma}^{0,\text{N}^{\text{k}}\text{LO}}$ denotes the bare cross-section prior to coupling-constant renormalisation.

2.3.2. Infrared divergences

Matrix elements are functions of Mandelstam variables

$$s_{ij} = (p_i + p_j)^2 = 4E_i E_j \sin^2 \frac{\theta_{ij}}{2}, \quad (2.3.25)$$

where i and j index massless (or effectively-massless) partons in the hard scattering, and θ_{ij} is the angle between them. If i and j are both identified particles in the final state, E_i and E_j must each be large enough for particles i and j to be separately detected by experimental calorimeters (which have an energy threshold for detection), and θ_{ij} must be large enough for them to be resolved separately, rather than identified as a single particle. These constraints prevent $s_{ij} \rightarrow 0$ for identified final-state particles.

Where these conditions are not satisfied, however, such as for i or j an unresolved final state particle, any inverse powers of s_{ij} will lead to divergences in the matrix element. Since they can arise for small E , they are called ‘infrared’ (IR) divergences. They can be traced back to Feynman diagrams with a propagator of momentum $p_i + p_j$, which contributes a divergent factor of this type when the momentum goes on-shell. Whether such a diagram exists (or its vertex factor is zero) depends on the partonic identity of i and j .

We have already seen from the splitting functions of eqs. (2.1.7) to (2.1.10) that there is a universal factorisation at the matrix-element level, in the limit in which two partons become collinear, exploited in the calculation of the DGLAP equations. There is a similar factorisation in the limit in which one parton becomes soft (where the reduced matrix element exists, and is not zero). This universal factorisation behaviour allows the factors to be extracted from the matrix elements for one process

and applied to those for others. This is the underlying principle for the method of antenna subtraction, which will be introduced in chapter 3.

These IR divergences also occur in loop diagrams. For example, the bubble diagram of eq. (2.3.1) can be seen from eq. (2.3.2) to contain factors of the form

$$\log^n \left(\frac{-p^2}{\mu^2} \right), \quad (2.3.26)$$

where we have included the renormalisation factor μ^ε from eq. (2.3.6). As a result, the integral will diverge for momenta that are small relative to the regularisation scale. In QCD, these virtual divergences are universal and depend only on the colour structure. In general, in a massless theory they arise in one-loop diagrams when [61]:

- an external on-shell particle is attached to two internal propagators, and when
- two on-shell external propagators exchange a virtual particle.

These correspond to collinearity of a loop propagator with the external momentum, and vanishing propagator momentum transfer respectively, so can be classified as collinear and soft singularities as in the real emission case.

This connection between real and virtual IR singularities will be seen to have a deep significance for perturbative calculations through the Kinoshita–Lee–Nauenberg (KLN) theorem in section 2.6. First we summarise the IR behaviour of each.

2.3.2.1. Virtual IR singularities

The universal IR structure of virtual corrections is known fully for QCD amplitudes corresponding to diagrams of up to two loops [62, 63], expressed in terms of Catani pole operators. These act on colour-space QCD amplitudes and isolate the principal part of the Laurent series.

For example, at one-loop,

$$\mathbf{P} \left[\left| \mathcal{M}_n^{(1)}(\varepsilon, \mu^2; \{p\}) \right\rangle \right] = \mathbf{I}^{(1)}(\varepsilon, \mu^2; \{p\}) \left| \mathcal{M}_n^{(0)}(\mu^2; \{p\}) \right\rangle \quad (2.3.27)$$

where the remainder is regular in ε and so finite in the $\varepsilon \rightarrow 0$ limit, and must be calculated separately for each process. The operator is given by

$$\mathbf{I}^{(1)}(\varepsilon, \mu^2; \{p\}) = \frac{e^{\varepsilon\gamma_E}}{2\Gamma(1-\varepsilon)} \sum_{i \in \{q, \bar{q}, g\}} \left[\frac{1}{\varepsilon^2} + \frac{\gamma_i}{C_i} \frac{1}{\varepsilon} \right] \sum_{j \neq i} \mathbf{T}_i \cdot \mathbf{T}_j \left(\eta_{ij} \frac{\mu^2}{s_{ij}} \right)^\varepsilon, \quad (2.3.28)$$

where i and j denote (QCD) partonic identities, the colour factors arise from the Casimirs of eq. (1.1.50) and the QCD beta-function eq. (2.3.15),

$$C_q = C_{\bar{q}} = C_F = \frac{N_c^2 - 1}{2N_c}, \quad C_g = C_A = N_c, \quad (2.3.29)$$

$$\gamma_q = \gamma_{\bar{q}} = \frac{3}{2}C_F, \quad \gamma_g = \beta_0, \quad (2.3.30)$$

the generalised colour-space operators $\mathbf{T}_i$ are the colour algebra generators in the representation of i (with an additional factor of -1 for initial-state quarks and final-state antiquarks), and the unitarity phase is given by

$$\eta_{ij} = \begin{cases} e^{-i\pi} & \text{if } i, j \text{ both incoming or outgoing} \\ 1 & \text{otherwise.} \end{cases} \quad (2.3.31)$$

For example, for processes with a colourless final-state X , such as diphoton production, this gives for $q\bar{q} \rightarrow X$ the divergent part of the one-loop matrix element,

$$\begin{aligned} \langle \mathcal{M}_n^{(0)} | \mathbf{I}^{(1)}(\varepsilon, \mu^2; \{1_q, 2_{\bar{q}}\}) | \mathcal{M}_n^{(0)} \rangle &= -C_F \frac{e^{\varepsilon\gamma_E}}{\Gamma(1-\varepsilon)} \left[\frac{1}{\varepsilon^2} + \frac{3}{2\varepsilon} \right] \\ &\times \left(e^{-i\pi} \frac{\mu^2}{s_{12}} \right)^\varepsilon \left| \mathcal{M}_n^{(0)}(1_q, 2_{\bar{q}}, X) \right|^2. \end{aligned} \quad (2.3.32)$$

The two-loop singularities can be expressed in a similar way [64]:

$$\begin{aligned} \text{P} \left[\left| \mathcal{M}_n^{(2)}(\varepsilon, \mu^2; \{p\}) \right| \right] &= \mathbf{I}^{(2)}(\varepsilon, \mu^2; \{p\}) \left| \mathcal{M}_n^{(0)}(\mu^2; \{p\}) \right| \\ &+ \mathbf{I}^{(1)}(\varepsilon, \mu^2; \{p\}) \left| \mathcal{M}_n^{(1)}(\varepsilon, \mu^2; \{p\}) \right|, \end{aligned} \quad (2.3.33)$$

where again the regular part of the Laurent series, including the finite remainder in the $\varepsilon \rightarrow 0$ limit, must be calculated separately for each process. The precise form of $\mathbf{I}^{(2)}$ can be found in [63].

2.3.2.2. Real IR singularities

Infrared singularities arising from unresolved outgoing particles also satisfy a universal factorisation relating higher-order amplitudes to simpler lower-order ones. The universal factorisation in the collinear limit has already been exploited to justify the DGLAP equations of section 2.1.2. Here we summarise the universal factorisation in the infrared limits.

In QED the soft-photon bremsstrahlung amplitude factorises neatly into the product of a soft ‘eikonal’ factor with a reduced matrix element [65, 66]

$$M_{n+m}^{(0)}(p_1, \dots, p_n; k_1, \dots, k_m) = \left[\prod_{i=1}^m S(p_1, \dots, p_n; k_i) \right] M_n^{(0)}(p_1, \dots, p_n), \quad (2.3.34)$$

but in QCD, the additional group-theory factor in the gluon vertex compared to the QED photon vertex substantially complicates the divergences of iterative emissions:

$$S_{\text{QED}}(p; k_\gamma) = e \frac{p \cdot \varepsilon(k)}{p \cdot k} \mapsto S_{\text{QCD}}(p; k_g) = g_s T_{ij}^a \frac{p \cdot \varepsilon(k)}{p \cdot k}. \quad (2.3.35)$$

The QCD factorisation in the unresolved limits only manifests itself when expressed in terms of colour-ordered amplitudes, in which the colour structures are factorised from the kinematics of the Feynman integrals. For example, in the case of a tree-level all-gluon amplitude [67, 68],

$$M_n^{(0)}(1, \dots, n) = g^{n-2} \sum_{\sigma \in S_n/Z_n} \text{Tr}[T^{a_{\sigma(1)}} \dots T^{a_{\sigma(n)}}] \mathcal{A}(\sigma(1), \sigma(2), \dots, \sigma(n)), \quad (2.3.36)$$

and in the case of m gluons emitted from a quark line,

$$M_{m+2}^{(0)}(q; 1, \dots, m; \bar{q}) = g^m \sum_{\sigma \in S_3} [T^{a_{\sigma(1)}} \dots T^{a_{\sigma(m)}}]_{ij} \mathcal{A}(q; \sigma(1), \dots, \sigma(m); \bar{q}), \quad (2.3.37)$$

where the $\mathcal{A}$ are called ‘colour-ordered subamplitudes’. When the amplitudes are squared (as in eq. (1.2.34)) we can apply the Fierz identity eq. (1.1.46),

$$T_{ij}^a T_{kl}^a = \frac{1}{2} \left(\delta_{il} \delta_{jk} - \frac{1}{N_c} \delta_{ij} \delta_{kl} \right) \quad (2.3.38)$$

to extract a series expansion in the inverse of the number of colours squared, $1/N_c^2$,

$$\left| M_n^{(0)} \right|^2 = \left(g^2 N_c \right)^{n-2} \left(N_c^2 - 1 \right) \left[\sum_{\sigma \in S_n/Z_n} M_n^{(0)}(\sigma(1), \dots, \sigma(n)) - \frac{1}{N_c^2} \tilde{M}_n^{(0)} + \dots \right] \quad (2.3.39)$$

where we denote sub-leading colour corrections, which are interference terms of different colour-ordered subamplitudes, with a corresponding number of tildes. This notation is useful, since these correspond to the interference of subamplitudes in which a gluon is abelian, charged under $U(1)$ rather than $SU(N_c)$, has no non-abelian couplings and hence simply factorises in these limits as a photon would.

Single unresolved limits

At tree-level, in the soft-gluon limit, the factorisation of a colour-ordered matrix element $M_{n+1}^{(0)}$ is QED-like:

$$M_{n+1}^{(0)}(\dots, i, j_g, k, \dots) \xrightarrow{p_j \rightarrow 0} 2 \frac{s_{ik}}{s_{ij}s_{jk}} M_n^{(0)}(\dots, i, k, \dots). \quad (2.3.40)$$

The soft-quark limit vanishes, as the corresponding reduced matrix element would violate quark-number conservation and so must be zero.

In the collinear limit (upon spin-averaging¹⁴), it factorises as

$$M_{n+1}^{(0)}(\dots, i, j, \dots) \xrightarrow{i \parallel j} \frac{P_i^K(\xi_i)}{s_{ij}} M_n^{(0)}(\dots, K, \dots) \quad (2.3.41)$$

where $p_i \rightarrow \xi_i p_K$, and P_i^K is one of the Altarelli-Parisi splitting functions of section 2.1.2.

Double unresolved limits

At NNLO we must consider the radiative corrections arising from the emission of two unresolved partons, and so two particles can simultaneously become soft, one

¹⁴The corresponding spin-dependent collinear factorisation formula also exists, but is only required when we do not intend to perform the full final-state phase-space integral, since the spin-dependent and spin-averaged versions differ only by azimuthal terms proportional to $\cos 2\phi$ which cancel in the phase-space integral over ϕ .

can become soft whilst the other becomes collinear to a hard parton, or two particles can become collinear to hard partons.

In each case the singularity structure depends on the colour-connectedness of the two partons. Where they are colour-disconnected the two limits are unrelated and can be taken separately, which gives factorisation identical to that in the NLO case, iterated. For example, in the double-collinear limit, in which the two unresolved partons are collinear with distinct hard partons, [64, 67]

$$M_{n+2}^{(0)}(\dots, i, j, \dots, k, l, \dots) \xrightarrow[p_k \parallel p_l]{p_i \parallel p_j} \left(\frac{P_i^I(\xi_i)}{s_{ij}} \right) \left(\frac{P_k^K(\xi_k)}{s_{kl}} \right) M_n^{(0)}(\dots, I, \dots, K, \dots), \quad (2.3.42)$$

and similarly for the soft-collinear limit, and the double-soft-gluon limit.

When the unresolved partons are colour-connected, new NNLO singular factors arise, which are more complicated functions of the invariants than those occurring at NLO. For instance, in the case of a quark-antiquark pair simultaneously going soft, [69]

$$M_{n+2}^{(0)}(\dots, i, j_q, k_{\bar{q}}, l \dots) \xrightarrow{p_j, p_k \rightarrow 0} S_{il}(j_q, k_{\bar{q}}) M_n^{(0)}(\dots, i, l, \dots), \quad (2.3.43)$$

where the universal soft-quark singularity function is [70]

$$S_{il}(j_q, k_{\bar{q}}) = \frac{2}{s_{jk}^2} \left(\frac{s_{il}s_{jk} - s_{ij}s_{kl} - s_{jl}s_{ik}}{(s_{ij} + s_{ik})(s_{jl} + s_{kl})} + \frac{s_{ij}s_{ik}}{(s_{ij} + s_{ik})^2} + \frac{s_{jl}s_{kl}}{(s_{jl} + s_{kl})^2} \right). \quad (2.3.44)$$

Similarly, for a pair of colour-connected soft gluons between two hard radiators, we see a similar factorisation

$$M_{n+2}^0(\dots, i, j_g, k_g, l \dots) \xrightarrow{p_j, p_k \rightarrow 0} S_{il}(j_g, k_g) M_n^0(\dots, i, l, \dots), \quad (2.3.45)$$

where the soft function $S_{il}(j_g, k_g)$ is [69]

$$S_{il}(j_g, k_g) = 2 \left[\frac{s_{il}^2}{s_{ij}s_{ijk}s_{jkl}s_{kl}} + \frac{s_{il}}{s_{jk}} \left(\frac{1}{s_{ij}s_{kl}} + \frac{1}{s_{ij}s_{jkl}} + \frac{1}{s_{ijk}s_{kl}} - \frac{4}{s_{ijk}s_{jkl}} \right) \right. \\ \left. + (1 - \varepsilon) \frac{1}{s_{jk}^2} \left(\frac{s_{ij}}{s_{ijk}} + \frac{s_{kl}}{s_{jkl}} - 1 \right)^2 \right]. \quad (2.3.46)$$

In the triple-collinear limit, where three colour-connected partons become simultaneously collinear,

$$M_{n+2}^{(0)}(\dots, i, j, k, \dots) \xrightarrow{p_i \parallel p_j \parallel p_k} P_{ijk}^L(\xi_i, \xi_j, \xi_k) M_n^{(0)}(\dots, L, \dots) \quad (2.3.47)$$

where L is the composite parton, and $P_{ijk}^L(\xi_i, \xi_j, \xi_k)$ is a triple-collinear splitting function dependent on the partonic nature of i, j, k and L , with longitudinal momentum fractions ξ_i, ξ_j, ξ_k respectively (with the constraint that they sum to 1). The five distinct triple-collinear splitting functions are relatively complicated functions of the kinematic variables and are given in [69].

Loop-level single unresolved limits

At NNLO, the real-virtual matrix elements can also have unresolved external partons. These matrix elements arise from the interference of a one-loop amplitude with a tree-level amplitude, as in eq. (2.2.8), and so the factorisation must be modified to include the contribution of loop diagrams in the limit:

$$M_{n+1}^{(1)}(\dots, i, j, k, \dots) \xrightarrow{p_j \text{ unresolved}} f_{ijk}^{(1)}(p_i, p_j, p_k) M_n^{(0)}(\dots, i, k, \dots) + f_{ijk}^{(0)}(p_i, p_j, p_k) M_n^{(1)}(\dots, i, k, \dots) \quad (2.3.48)$$

where the one-loop soft- and splitting-functions are given in [71].

2.3.3. Infrared safety

We have seen that there is no meaningful physical distinction that can be drawn between n -particle final-states, and $(n + m)$ -particle final states containing an additional m partons that cannot be individually resolved. We cannot, for instance, meaningfully calculate a final-state with a defined number of partons, since that would amount to a veto on the soft or collinear splittings which we cannot identify to be distinct final states. To properly define observable quantities, we must be careful not to artificially distinguish between such ‘degenerate’ final-states.

Such observables are called ‘infrared-safe’, and must satisfy, for any indexing of the final-state momenta, the ‘Sterman–Weinberg’ criteria [72],

$$O_{n+1}(\{p_i\}_{i=1}^n \cup \{p_{n+1}\}) \longrightarrow O_n(\{p_1, \dots, p_i + p_{n+1}, \dots, p_n\}) \quad (2.3.49)$$

in the limit in which p_i becomes collinear with p_{n+1} , or p_{n+1} becomes soft.

The perturbative calculation is only well-defined for observables meeting this criterion. This leads to the concept of a ‘jet’, which will be described in further detail in section 2.5. The fact that the calculation is then well-defined then follows from an extension of the KLN theorem [73, 74], which states that infrared divergences in massless gauge theories cancel in the sum over all degenerate initial- and final-states. This is itself an extension of the Bloch–Nordsieck theorem for QED [75], which proves a similar result arising from the sum over degenerate final-states alone.

In practice, the cancellation occurs in the Laurent series computed through dimensional regularisation. The infrared structure of loop diagrams is calculated as an explicit Laurent series in ε , whilst the singularities associated with unresolved additional partons only manifest themselves as ε -poles after the integration over phase-space in d -dimensions.

This suggests an alternative approach. The KLN theorem, and the eventual cancellation of these ε -poles it guarantees, can be exploited to define local counterterms that mimic the divergent structure of the matrix elements in the unresolved limits. For example, if we suppose that at NLO we can define $d\hat{\sigma}^S$ to mimic the divergent behaviour of $d\hat{\sigma}^R$ in the relevant regions of phase-space, the integral

$$\int d\Phi_{n+1} \left[\frac{d\hat{\sigma}^R}{d\Phi_{n+1}} - \frac{d\hat{\sigma}^S}{d\Phi_{n+1}} \right] \quad (2.3.50)$$

is everywhere finite, and so can safely be computed numerically in $d = 4$ dimensions. This is the underlying principle for the subtraction method discussed in section 2.6, and for the specific approach to computing the subtraction counterterms used in the remainder of this work, antenna subtraction, described in chapter 3.

2.3.4. PDF renormalisation and mass factorisation

The KLN theorem guarantees the cancellation of infrared singularities in the inclusive sum over degenerate initial- and final-states, but within the parton model we sum only over degenerate final-states and impose definite initial states through the factorisation with specific PDFs corresponding to a certain partonic initial-state. We therefore have collinear divergences remaining in the partonic cross-sections associated with these initial-state singularities.

For massless particles, numerically these can be absorbed into a redefinition of the PDFs¹⁵ by the convolution

$$f_a(\xi; \mu_F^2) = \sum_b \int_0^1 dx dy \Gamma_{ab}(x, \mu_F) f_b^0(y) \delta(\xi - xy) \quad (2.3.51)$$

$$=: \sum_b [\Gamma_{ab}(\mu_F) \otimes f_b^0] \quad (2.3.52)$$

where the factorisation kernels Γ_{ab} are renormalisation-scheme dependent. This is the origin of the μ_F -dependence that was assumed on physical grounds in section 2.1, and which is absent from the ‘bare’ PDFs $\{f_a^0(\xi)\}$.

The factorisation kernels $\Gamma_{ab}(x)$ have perturbative expansion

$$\begin{aligned} \Gamma_{ab}(x, \mu) = \delta_{ab} \delta(1-x) + \left(\frac{\alpha_s(\mu)}{2\pi} \right) \Gamma_{ab}^{(1)}(x) \\ + \left(\frac{\alpha_s(\mu)}{2\pi} \right)^2 \Gamma_{ab}^{(2)}(x) + \mathcal{O}(\alpha_s(\mu)^3), \end{aligned} \quad (2.3.53)$$

where in the $\overline{\text{MS}}$ scheme [76, 77],

$$\Gamma_{ab}^{(1)}(x) = -\frac{1}{\varepsilon} P_a^{b(0)}(x) \quad (2.3.54)$$

$$\Gamma_{ab}^{(2)}(x) = \frac{1}{2\varepsilon^2} \left(\sum_c [P_a^{c(0)} \otimes P_c^{b(0)}](x) + 2\beta_0 P_a^{b(1)}(x) \right) - \frac{1}{2\varepsilon} P_a^{b(1)}(x). \quad (2.3.55)$$

These contributions can be expanded to establish the order at which they enter the cross-section. For example, at NLO the contribution from this ‘mass-factorisation’

¹⁵The full picture is slightly subtler, involving an interplay between IR and UV divergences, and is outlined in chapter 9 of [44].

counterterm is

$$\begin{aligned} d\hat{\sigma}_{ab}^{\text{MF NLO}} = -C_\varepsilon \sum_c \int \frac{dx_a}{x_a} \frac{dx_b}{x_b} \left[\delta(1-x_a) d\hat{\sigma}_{ac}^{\text{B}} \Gamma_{cb}^1(x_b) \right. \\ \left. + \Gamma_{ca}^1(x_a) d\hat{\sigma}_{cb}^{\text{B}} \delta(1-x_b) \right]. \end{aligned} \quad (2.3.56)$$

A detailed discussion of the mass-factorisation contribution at NNLO is deferred until section 3.5.3.4, when it will be applied to diphoton production within the antenna subtraction formalism.

2.4. Scale dependence of QCD cross-sections

In order to make concrete predictions using the parton model, we must choose a renormalisation scale μ_{R} and a factorisation scale μ_{F} . Whilst at all orders, the predictions must be independent of these unphysical variables, at a finite order the dependence remains. This can be derived analytically from the renormalisation group equations, and provides a powerful check on the accuracy of numerical results. This will be summarised in section 2.4.1 and applied to verify our calculation in appendix A.3.2.

Because the missing corrections are also functions of the scales, if at some sufficiently high order of the perturbative expansion the dependence vanishes, the scale-dependence of the missing corrections must be the inverse of that of the calculable terms, yielding a constant independent of the chosen scale.

This is often used in practice to provide an estimation of the size of the missing-higher-orders. It is expected that the missing scale-dependent terms will be dominated by logarithmic functions of ratios of scales,

$$L_{12} := \log \frac{\mu_1}{\mu_2}. \quad (2.4.1)$$

If this is the case, then the corrections will be largest where the scales μ_1 and μ_2 are of different orders of magnitude, and small where they are of comparable orders of magnitude. Choosing a dynamical scale variable which intuitively matches the

relevant physical scale of the process aims to ensure that the largest ratio is bounded, which would be impossible to guarantee with a constant scale.

2.4.1. Analytical dependence

The scale dependence of a QCD cross-section arises separately from:

1. the factorisation-scale dependence of the PDFs, through their renormalisation-group equations (the DGLAP equations, eq. (2.1.6)), and
2. the renormalisation of the strong coupling, through its renormalisation-group equation (the QCD beta function eq. (2.3.14)).

The dependence of the cross-section on the renormalisation and factorisation scales at a given order can therefore be extracted from the iterative solutions to these two equations.

Here we restrict to the renormalisation-scale dependence of the cross-section, effectively always evaluating PDFs at a common factorisation scale. Iteratively applying eq. (2.3.17) gives

$$\begin{aligned} \alpha_s(Q) = \alpha_s(\mu) & \left[1 + \left(\frac{\alpha_s(\mu)}{4\pi} \right) \beta_0 \log \frac{\mu^2}{Q^2} \right. \\ & + \left(\frac{\alpha_s(\mu)}{4\pi} \right)^2 \left\{ \left(\beta_0 \log \frac{\mu^2}{Q^2} \right)^2 + \beta_1 \log \frac{\mu^2}{Q^2} \right\} \\ & \left. + \mathcal{O}(\alpha_s^3) \right]. \end{aligned} \quad (2.4.2)$$

Inserting this into the perturbative expansion of the cross-section, eq. (2.2.1), gives

$$\begin{aligned} d\hat{\sigma}(\mu_R, \alpha_s(\mu_R)) &= \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^m d\hat{\sigma}_0^{\text{LO}} \\ &+ \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{m+1} \left[d\hat{\sigma}_0^{\text{NLO}} + \frac{1}{2} m \beta_0 \log \frac{\mu_R^2}{\mu_0^2} d\hat{\sigma}_0^{\text{LO}} \right] \\ &+ \left(\frac{\alpha_s(\mu_R)}{2\pi} \right)^{m+2} \left[d\hat{\sigma}_0^{\text{NNLO}} + \frac{1}{2} (m+1) \log \frac{\mu_R^2}{\mu_0^2} d\hat{\sigma}_0^{\text{NLO}} \right] \end{aligned}$$

$$\begin{aligned}
& + \frac{1}{4} m \left(\beta_1 + \frac{1}{2} (m+1) \beta_0 \log \frac{\mu_R^2}{\mu_0^2} \right) \log \frac{\mu_R^2}{\mu_0^2} d\hat{\sigma}_0^{\text{LO}} \Big] \\
& + \mathcal{O}(\alpha_s^{m+3}), \tag{2.4.3}
\end{aligned}$$

where at each order

$$d\hat{\sigma}_0^{\text{N}^k\text{LO}} := d\hat{\sigma}^{\text{N}^k\text{LO}}(\mu_0, \alpha_s(\mu_0)). \tag{2.4.4}$$

2.4.2. Estimation of missing higher-orders

The conventional scale variation procedure involves the calculation of fixed-order cross-sections and differential cross-sections at factorisation and renormalisation scales scaled up and down from the chosen central scale μ_0 by a constant factor r , and taking the resulting envelope. Typically the factor $r = 2$ is chosen, and the ‘scale uncertainty’ band defined by scale-variation is defined as the envelope $[\sigma_-, \sigma_+]$ around the central-scale value σ , with

$$\sigma_- = \min_{k_1, k_2 \in S_r} \{ \sigma(\mu_F = k_1 \mu_0, \mu_R = k_2 \mu_0) \} \tag{2.4.5}$$

$$\sigma_+ = \max_{k_1, k_2 \in S_r} \{ \sigma(\mu_F = k_1 \mu_0, \mu_R = k_2 \mu_0) \} \tag{2.4.6}$$

where the conventional 7-point set of factors is

$$S_r = \left\{ \left(\frac{1}{r}, \frac{1}{r} \right), \left(\frac{1}{r}, 1 \right), \left(1, \frac{1}{r} \right), (1, 1), (1, r), (r, 1), (r, r) \right\}. \tag{2.4.7}$$

This attempts to estimate the magnitude of (unknown) missing higher-order corrections to fixed-order calculations through the heuristic argument [78]

$$\left| \hat{\sigma}_{\mu_R=rQ}^{(k)} - \hat{\sigma}_{\mu_R=Q/r}^{(k)} \right| \approx \left| \frac{d\hat{\sigma}^{(k)}}{d \ln \mu_R^2} \right|_{\mu_R=Q} (4 \ln r) \tag{2.4.8}$$

$$\approx 4 \alpha_s(Q)^{k+1} k \beta_0 \ln r |c_k|, \tag{2.4.9}$$

$$\gtrsim \alpha_s(Q)^{k+1} |c_{k+1}| \approx \sum_{n=k+1}^{\infty} c_n \alpha_s^n, \tag{2.4.10}$$

where the fixed-order truncation of the perturbative series is the partial sum

$$\hat{\sigma}^{(k)}(Q) = \sum_{n=0}^k c_n(Q, \mu_R) \alpha_s(\mu_R)^n \quad (2.4.11)$$

and we assume that the coefficients c_n are all of the same order of magnitude.

This leads, in principle, to each successive order of the calculation giving central predictions lying within the scale uncertainty bands of the previous order, with a narrower uncertainty band in which higher orders are expected to lie, illustrating convergence of the perturbative series.

This ‘scale uncertainty’, reflecting the size of the unknown terms in the perturbative series, is really a truncation uncertainty, and is not a conventional uncertainty as no probabilistic interpretation of ‘degree of belief’ can be assigned to it (though recent attempts have been made to assign one through a Bayesian framework in [78–80]). It is instead a heuristic that is used in practice to interpret the quality of a prediction and its likely predictive power.

This is problematic as a procedure if there are *a priori* reasonable choices of scales which differ by a factor greater than r . Often, reasonable choices of dynamic scales will differ in some kinematic regions by substantially more than a factor of r . Typically such regions arise when the asymptotic behaviour of the two scales differs along some axis, and the fiducial cuts define a sufficiently large region of phase-space to allow the asymptotic behaviour to generate large ratios.

As a result, the results in such regions are especially sensitive to our choice of functional form for the central scale, and naïve scale variation of factors of r about the central scale will underestimate the true envelope of scale uncertainty in these regions, by generating uncertainty bands of width proportional to $\ln r$ rather than the ratio of scales

$$\ln \frac{\mu'}{\mu} > \ln r. \quad (2.4.12)$$

2.4.3. Scale-setting techniques

2.4.3.1. New scales for old

Many scales arise naturally in a high-energy scattering. *A priori*, the scales at which the two hadrons are probed are not necessarily the same, which would suggest the use and separate variation of two, different, factorisation scales. Every pair of final-state particles i and j defines its own four-momentum invariant, and due to the identity (for massless particles)

$$s_{ij} = 2p_{\text{T}}^i p_{\text{T}}^j \left(\cosh \Delta y_{ij} - \cos \Delta \phi_{ij} \right), \quad (2.4.13)$$

these can range over several orders of magnitude even between pairs of particles with identical projections onto the transverse plane, due to rapidity separation alone.

It is not obvious which of these to choose for a perturbative calculation. However, intrinsically multi-scale problems can easily be converted to a single-scale problem by taking an average. Applying eq. (2.3.18),

$$\prod_{i=1}^n \alpha_{\text{s}}(\mu_i) = \alpha_{\text{s}}(\mu_0)^n \prod_{i=1}^n \left[1 + \beta_0 \frac{\alpha_{\text{s}}(\mu_0)}{4\pi} \log \left(\frac{\mu_0^2}{\mu_i^2} \right) + \mathcal{O}(\alpha_{\text{s}}(\mu_0)^2) \right], \quad (2.4.14)$$

$$= \alpha_{\text{s}}(\mu_0)^n \left[1 + \beta_0 \frac{\alpha_{\text{s}}(\mu_0)}{4\pi} \log \left(\frac{\mu_0^{2n}}{\prod_{i=1}^n \mu_i^2} \right) + \mathcal{O}(\alpha_{\text{s}}(\mu_0)^2) \right]. \quad (2.4.15)$$

As a result, the choice of the geometric mean,

$$\mu_0 = \left(\prod_{i=1}^n \mu_i \right)^{\frac{1}{n}} \quad (2.4.16)$$

ensures that the difference between the choice of the individual scales μ_i for each factor of α_{s} , and the choice of their geometric mean for all of them, is formally of $\mathcal{O}(\alpha_{\text{s}}^2)$, whilst for any other choice of μ_0 it is formally $\mathcal{O}(\alpha_{\text{s}}^1)$.

2.4.3.2. Strategies for scale-setting

Given the residual dependence of the prediction on the choice of scales, a number of strategies have been developed to give the ‘optimal’ scale. This remains an

open problem, and none of these approaches has been widely adopted for NNLO calculations. As a general rule, the properties they are conjectured to confer on predictions made with their proposed choice of scale could only be verified to a given order either by a formal proof, or a calculation to a higher order demonstrating the desired characteristics. As such, their conjectured advantages for phenomenology at the highest calculable order, although motivated, are not guaranteed to hold.

Since the factorisation scale dependence of the cross-section arises through the PDFs, which are extracted from fits of predictions to data, there is some circularity in the choice of factorisation scale: a PDF set produced with fits of predictions using one set of scales is not guaranteed to produce as good a fit to the same data when used with another. Determining the resulting uncertainty and propagating it to errors in predictions made with the PDFs in a theoretically valid way is an open problem. These uncertainties can only possibly be accounted for in the PDF fitting stage, and PDF fitting collaborations are only beginning to consider how to include them systematically in their fits [81–84]. We therefore focus on the choice of the renormalisation scale, with the proviso that in practice it is customary to use the same choice for the central factorisation scale too.

Fastest Apparent Convergence

The principle of Fastest Apparent Convergence was introduced in [85]. It sets the renormalisation scale at the value which causes the NLO correction to vanish, with the motivation that the MHOU's are expected to contain higher powers of logarithms of the same ratios, which might therefore also vanish.

Principle of Minimal Sensitivity

The Principle of Minimal Sensitivity [86–89] *imposes* the local invariance of the truncated perturbative series under changes in the renormalisation scale, and solves for the value of the scale at which it occurs (if it exists, which is not guaranteed).

Brodsky-Lepage-Mackenzie

The Brodsky-Lepage-Mackenzie (BLM) criterion [90] determines a separate ‘optimal’ renormalisation scale to be used for each successive order of perturbative corrections. At each order, it chooses a scale that aims to absorb the terms proportional to the QCD β -function coefficients into the running coupling, using the n_f terms (from the quark vacuum polarisation) to identify them. However, it is not well-defined at higher orders, as the powers of n_f no longer uniquely define linear combinations of (products of) β -function coefficients [91–93].

Principle of Maximum Conformality

The Principle of Maximum Conformality [92, 94–96] extends the motivation behind the BLM criterion to absorbing all terms proportional to the β -function into the running coupling. The remaining terms in the series are then identical to those that would arise for a conformal theory, in which the β -function is identically zero and all coefficients vanish, and the coupling is independent of the scale.

2.5. Jets and hadrons

We saw in section 2.1 that colour confinement led to a distinction between the colour-singlet particles that can be isolated and hence observed, and the constituent (colour-charged) fields of QCD, which behave as particles inside high-energy collisions. This observation led to the parton model, in which the physics of the partons within each hadron was separated from the high-energy scattering of the interacting partons.

The same phenomenon, in reverse, leads to the rapid recombination of coloured final-state partons and their radiated particles into an energetic spray of collimated hadrons, which are detected by hadronic calorimeters in collider experiments.

Jet algorithms provide rules for grouping these particles into ‘jets’, specifying which particles to combine and the rules for combining them, assigning a mass and four-momentum to the composite object.

Once defined, these algorithms can be applied to any collection of objects with the required properties. These could be partons in the final state of the hard scattering, those produced after parton showering, theoretical models of hadrons, measured energy depositions in calorimeter cells, or particle objects reconstructed from calorimeter tracks. Detector effects can be modelled and ‘unfolded’ to correct for their effect on data, but the hadronisation process is intrinsically non-perturbative and cannot be modelled theoretically at an accuracy comparable to the perturbative parts of the calculation.

Ideally, we would define a jet algorithm that would be transparent to hadronisation, and would, whatever the details of a model of hadronisation turn out to be, recombine the resulting hadrons into the same jets as their parent partons. In fact, we go further, and hope to define a single jet algorithm that defines the same jets no matter which level of the theoretical calculation (hard scattering, parton showering, hadronisation) or experimental measurement (calorimeter cells or reconstructed particle) it is applied to.

This is shown schematically in fig. 2.5. The validity of commuting the steps in the diagram determines the validity of comparing the theoretical parton-level jets of predictions with experimental jets reconstructed from calorimeter data at colliders.

2.5.1. Jet algorithms

There are two main categories of jet algorithm: cone-based, and sequential reclustering. Broadly, these reflect a top-down, and a bottom-up approach, respectively: cone-based algorithms identify a detector region around a jet axis and cluster together all particles within the region, whilst sequential algorithms recursively recombine the closest objects in an event (according to some metric) until a termination condition is met.




Figure 2.5: Levels of a calculation or measurement at which a jet algorithm can be applied to give a jet observable. The levels on the top line are purely theoretical, whilst those on the bottom are purely experimental. The goal is to reconcile these theoretical and experimental jets to enable the comparison of parton-level predictions with experimental collider data.

2.5.1.1. Cone algorithms

Cone algorithms [72,97–99] rely on the idea that neither soft, nor collinear, splittings much alter the underlying direction of energy flow. As a result, they assume that a hard parton’s final-state radiation will predominantly be emitted in a conical region around it, and so impose circular boundaries on the detector variables (η, ϕ) around some chosen axis, of radius

$$\Delta R = \sqrt{\Delta\phi^2 + \Delta\eta^2}. \quad (2.5.1)$$

All particles (or calorimeter cells) lying within the boundary are combined together to form a jet, resulting in jets of identical shape with sharply-defined geometric edges.

Early cone algorithms applied at colliders were not infrared-safe [100], relying, for example, on the hardest parton to define an initial jet axis. If this parton were to split into two collinear softer partons, a different axis could be chosen, resulting in a final-state with different identified jets.

The only infrared-safe cone algorithms must consider all possible cones rather

than choosing one to use as a seed. One example is enumerating all subsets of particles and establishing for each whether it corresponds to a stable cone (i.e. a circle around the axis of the subset's total momentum contains exactly the particles in the subset). There are $\mathcal{P}(\{i\}) = 2^N$ such subsets, and so this algorithm, naïvely implemented, is an $\mathcal{O}(N2^N)$ operation, which is prohibitive for experimental applications where N can be $\mathcal{O}(100)$. A polynomial-time solution to this problem, yielding an algorithm that identifies all such stable cones, SISCON, was developed in 2007 [99].

2.5.1.2. Sequential recombination

The favoured jet construction algorithms for use at the LHC are all sequential recombination algorithms, from a family collectively called ‘generalised k_T algorithms’.

These begin by defining a distance metric between two particles i and j ,

$$d_{ij} := \min \left\{ p_T^{i\ 2p}, p_T^{j\ 2p} \right\} \frac{\Delta R_{ij}^2}{R^2}, \quad (2.5.2)$$

and between a particle i and the beam B ,

$$d_{iB} := p_T^{i\ 2p}. \quad (2.5.3)$$

The parameter p acts to distort distances according to the transverse momentum of the particles involved, with either softer particles, or harder particles, being considered as closer together than they are purely geometrically, for $p < 0$ and $p > 0$ respectively. The special values $p = 1, 0, -1$ are given the names ‘ k_T ’ [101] (or ‘Durham’ [102]), ‘Cambridge/Aachen’ [103, 104], and ‘anti- k_T ’ [105] respectively, with anti- k_T typically used for jet studies at the LHC.

These algorithms differ only in their choice of p in the distance measure, and otherwise proceed identically. At each step, the shortest distance

$$\min_{i,j} \{d_{ij}, d_{iB}\} \quad (2.5.4)$$

is found. If the minimum is d_{iB} for some i , particle i is closer to the beam than any other particle, so is labelled a jet and removed from the list. If the minimum is d_{ij}




Figure 2.6: The different sequential-clustering jet algorithms applied to partons in the (η, ϕ) -plane with the same radius and transverse energy threshold. Each point corresponds to a parton, of p_T proportional to its radius. The coloured patches show the extent of the resulting jets when the partons are clustered with the anti- k_T , Cambridge-Aachen, and k_T algorithms respectively. In this case we see that all the algorithms identify four jets, but with slightly different properties.

for some i and j , they are combined into a single protojet (ij) with four-momentum $p_i + p_j$ and the set of distances is recalculated. The process terminates when no particles remain.

The different distance measures lead to very different phenomenological properties. The k_T -algorithm clusters soft particles preferentially, and so builds jets that are irregular and whose shape depends substantially on the soft radiation within the event. The anti- k_T algorithm clusters the hardest particles preferentially, identifying the cores of jets before filling them in with softer particles. In practice, this leads to regularly-shaped circular jets, making it an attractive choice. An example illustrating the properties of these alternatives is shown in fig. 2.6.

2.5.2. Measurement function formalism

At the level of the partonic cross-section, the imposition of a jet requirement on a final state can be incorporated into eq. (2.1.3) at the level of the integrand through a generalised jet function $J_m^n(\Phi_n)$, asserting that m jets have been found in the n -particle final-state. This is a Heaviside theta function that sets the integrand to

zero in the regions of phase-space where the jet requirements are not satisfied:

$$J_m^n(\Phi_n) = \begin{cases} 0 & < m \text{ jets in final-state} \\ 1 & \geq m \text{ jets in final-state.} \end{cases} \quad (2.5.5)$$

The specific form of J_m^n will depend on the jet definition chosen. Clearly if the jet definition is infrared-safe, the infrared singularities in the final-state discussed in section 2.3.2.2 can only occur when the number of identified jets is less than the number of partons; where they match, none of the particles is unresolved.

2.5.3. Jet substructure

The differences between the sequential-clustering jet algorithms can be used to reveal further information about events. As can be seen from fig. 2.6, the simple clustering of partons into jets removes information that might be valuable about the distribution and energies of the jet's constituents.

Jet substructure techniques [106–109] have been developed to look inside jets, allowing them to be further categorised. For instance, jets arising from hadronic decays typically have multiple hard ‘cores’, from each of the several decay products, whilst jets arising from the fragmentation of a single hard parton are dominated by the radiation of soft gluons, and so have a single core.

We illustrate the potential by focusing on a single example of these, soft-drop declustering [110]. Soft-drop aims to remove wide-angle soft radiation from a jet to reduce any contamination from background hadronic radiation from processes other than the hard scattering of interest.

A jet algorithm is first applied to the event, with some radius R to define jets. The constituents of each jet j are then reclustered using a different algorithm to form a tree of pairwise clusterings. The tree is then ‘pruned’. If j is a singleton (and so has a trivial clustering tree), it can either be discarded (leaving only composite objects as jets, ‘tagging’), or returned as a jet with no substructure (‘grooming’). If

it is composite, the final pair of proto-subjets is tested against the condition

$$\frac{p_T^2}{p_T^1 + p_T^2} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R} \right)^\beta, \quad (2.5.6)$$

where the jets are labelled by descending p_T so that $p_T^1 \geq p_T^2$, and z_{cut} and β are new parameters for the soft-drop algorithm.

If the condition is met, then j passes the declustering test, and is allowed to remain as an identified jet unchanged. If it is failed, the softer jet j_2 is discarded, so j_1 becomes j , and the highest-level clustering within j_1 is tested against eq. (2.5.6); if it passes it becomes the identified jet; if it fails, its harder proto-subjet j_{11} becomes j , and the process is iterated.

The name ‘soft drop’ arises because through eq. (2.5.6) we are requiring that the softer constituent subjet within a jet not be too soft relative to the jet (with the threshold a function of separation for $\beta \neq 0$), and ‘dropping’ it from the jet if it is. Because of the form of the threshold function, for $\beta > 0$ the criterion always permits (sufficiently) collinear radiation, which makes it IR unsafe in ‘tagging’ mode (as a collinear splitting changes the outcome).

We will see in chapter 4 how these ideas can be adapted to resolve photonic final-states.

2.6. The cancellation of IR divergences

The KLN theorem of section 2.3.3 guarantees the eventual cancellation of the ε -poles arising in loop diagrams after dimensional regularisation, against those arising from phase-space integration over regions of phase-space in which one or more partons are unresolved. Schematically, this cancellation takes the form

$$\underbrace{\frac{1}{\varepsilon}}_{\text{virtual pole}} + \underbrace{\left(\mu^2\right)^\varepsilon \int_0^{\lambda^2} \frac{dk^2}{\left(k^2\right)^{1+\varepsilon}}}_{\text{real divergence}} = \frac{1}{\varepsilon} - \frac{1}{\varepsilon} \left(\frac{\mu^2}{\lambda^2} \right)^\varepsilon \quad (2.6.1)$$

$$= \log \left(\frac{\lambda^2}{\mu^2} \right) + \mathcal{O}(\varepsilon) \quad (2.6.2)$$

where we have assumed $\varepsilon < 0$ to capture the IR divergence.¹⁶

Needing to integrate a divergent integrand in $d = 4 - 2\varepsilon$ dimensions prevents the application of numerical integration, since computer representations of real numbers do not permit the representation of arbitrarily large numbers and so would overflow rather than give a correct result. This problem arises because the cancellation of divergences occurs after integration, and not before.

To solve it, various methods have been proposed that allow the cancellation of divergences to occur before integration, at the integrand level. Here we outline the principles behind two of the common choices, slicing and subtraction, applied to a simple one-dimensional toy model of a divergent integral, before summarising the approaches taken to apply this to phase-space integrals of the form required for real-emission diagrams.

2.6.1. One-dimensional toy model

As a simple example, we consider [111]

$$I(\varepsilon) = \int_0^1 \frac{F(x)}{x} x^\varepsilon dx - \frac{1}{\varepsilon} F(0), \quad (2.6.3)$$

where F is some analytic function of x , bounded on $[0, 1]$ and potentially complicated enough to render the integral intractable. The x^ε factor regularises the pole in x at 0 and allows the integral to converge, provided that $\varepsilon > 0$. For $\varepsilon = 0$ the integral diverges, and so we will define $I(0)$ instead by the value taken in the limit $\varepsilon \rightarrow 0$. The ε -pole in the integral is exactly cancelled by the second term,¹⁷ rendering I analytic as a function of ε and so, in the limit, $I(0)$ is finite and only the ε^0 term will survive.

¹⁶Strictly, $\varepsilon \in \mathbf{C}$ and so such statements should be about $\text{Re}[\varepsilon]$, but we avoid this subtlety for simplicity of exposition.

¹⁷This can be seen by expanding $F(x)$ about $x = 1$ and identifying the integrals in each term of the power series as the Beta functions $B(-\varepsilon, n + 1)$.

2.6.1.1. Slicing

The idea of ('phase-space') slicing is to decompose the region of integration to isolate regions in which the integrand is singular, which will be integrated analytically, from regions in which it is well-behaved, which can be integrated numerically. In this case, the singular region is the neighbourhood of $x = 0$, and so we choose $\delta \ll 1$ for which we assume $F(x)$ can be approximated by $F(0)$ on $[0, \delta]$. Thus,

$$I(\varepsilon) \approx F(0) \int_0^\delta \frac{1}{x^{1-\varepsilon}} dx + \int_\delta^1 \frac{F(x)}{x^{1-\varepsilon}} dx - \frac{1}{\varepsilon} F(0) \quad (2.6.4)$$

$$= \frac{\delta^\varepsilon}{\varepsilon} F(0) + \int_\delta^1 \frac{F(x)}{x^{1-\varepsilon}} dx - \frac{1}{\varepsilon} F(0) \quad (2.6.5)$$

$$= \frac{1 + \varepsilon \log \delta + \mathcal{O}(\varepsilon^2)}{\varepsilon} F(0) + \int_\delta^1 \frac{F(x)}{x^{1-\varepsilon}} dx - \frac{1}{\varepsilon} F(0) \quad (2.6.6)$$

$$\xrightarrow{\varepsilon \rightarrow 0} F(0) \log \delta + \int_\delta^1 \frac{F(x)}{x} dx \quad (2.6.7)$$

The first term is easily calculable, and the second can be computed by numerical integration as the integrand is everywhere finite.

2.6.1.2. Subtraction

In contrast with 'slicing', the aim of subtraction methods is to construct a local counterterm that renders the integrand tractable to numerical quadrature (by Monte Carlo or otherwise).

In this case, we can mimic the singularity of the integrand at $x = 0$ with the counterterm

$$S(x) = \frac{F(0)}{x} x^\varepsilon, \quad (2.6.8)$$

which we can integrate analytically:

$$\int_0^1 S(x) dx = \frac{1}{\varepsilon} F(0) \quad (2.6.9)$$

and so

$$I(\varepsilon) = \int_0^1 \left[\frac{F(x)}{x} x^\varepsilon - S(x) \right] dx + \left[\int_0^1 S(x) dx - \frac{1}{\varepsilon} F(0) \right] \quad (2.6.10)$$

$$= \int_0^1 [F(x) - F(0)] \frac{x^\varepsilon}{x} dx, \quad (2.6.11)$$

$$\xrightarrow{\varepsilon \rightarrow 0} \int_0^1 \frac{1}{x} [F(x) - F(0)] dx. \quad (2.6.12)$$

where the cancellation of the second square bracket is serendipitous. This integrand is now finite, and so can be performed through numerical quadrature (including by Monte Carlo integration in higher dimensions). In practice, floating-point arithmetic necessitates the introduction of a technical cut-off parameter t to replace the lower limit of integration, as the loss of floating-point precision in the divergent limit makes the integrand numerically unstable, due to imperfect numerical cancellations of large numbers. By the construction of the subtraction term, we expect the missing region of integration to contribute approximately $tF'(0)$ to the final integral, and so can choose a t sufficiently small to render the error negligible.

Importantly, the choice of S made in eq. (2.6.8) was not unique. Any function with the same asymptotic behaviour in the divergent $x \rightarrow 0$ limit will regulate the singularity of the integrand. We can exploit this to choose a counterterm that is convenient to integrate analytically.

2.6.2. Application to phase-space integrals

The application of the above principles to phase-space integrals requires some adaptation.

Slicing closest to the schematic description above has typically been employed for NLO calculations of jet cross-sections [112, 113]. The real-emission phase-space is partitioned into regions using thresholds of kinematic invariants, chosen to distinguish the resolved from the unresolved regions. Where the emissions are guaranteed to be resolved, the exact matrix elements are finite, and numerical integration is unproblematic. In the regions where they may be unresolved, the matrix elements are approximated by the universal factorisation in the corresponding limit, which can be integrated analytically over the soft or collinear phase-space, leaving the reduced matrix element to be integrated numerically over the remaining dimensions.

As in the toy model, this introduces a dependence upon the threshold, which must be small enough to minimise the induced error from this approximation, but large enough that the resolved contribution remains finite and stable.

2.6.2.1. q_T - and N -jettiness subtraction

A related method to slicing, called q_T -subtraction [114, 115], has been used for the production of colourless final-states. This uses the transverse-momentum of the final state as a slicing variable:

$$\hat{\sigma}_{ab \rightarrow X}^{\text{N}^k\text{LO}} = \int_0^\infty dq_T \frac{d\hat{\sigma}_{ab \rightarrow X}}{dq_T} \quad (2.6.13)$$

$$= \int_0^{q_T^{\text{cut}}} dq_T \frac{d\hat{\sigma}_{ab \rightarrow X}}{dq_T} + \int_{q_T^{\text{cut}}}^\infty dq_T \frac{d\hat{\sigma}_{ab \rightarrow X}}{dq_T} \quad (2.6.14)$$

The Born kinematics must have $q_T = 0$, by momentum conservation with the colliding hadrons, so such contributions are proportional to $\delta(q_T)$. As a result, the parts of the calculation which inhabit the Born kinematics (Born, V, VV, etc.) contribute only to the first term, leaving in the second term only the contributions to $X + j$ at a lower order. That is, for $q_T > 0$,

$$\frac{d\hat{\sigma}_{ab \rightarrow X}^{\text{N}^k\text{LO}}}{dq_T}(q_T) = \frac{d\hat{\sigma}_{ab \rightarrow X+j}^{\text{N}^{k-1}\text{LO}}}{dq_T}(q_T), \quad (2.6.15)$$

and so

$$\hat{\sigma}_{ab \rightarrow X}^{\text{N}^k\text{LO}} = \int_0^{q_T^{\text{cut}}} dq_T \frac{d\hat{\sigma}_{ab \rightarrow X}}{dq_T} + \int_{q_T^{\text{cut}}}^\infty dq_T \frac{d\hat{\sigma}_{ab \rightarrow X+j}^{\text{N}^{k-1}\text{LO}}}{dq_T} \quad (2.6.16)$$

As a result, any method that can be used to handle the (simpler) real IR cancellations at a lower order can be extended one order higher. The small- q_T integrand still diverges $\sim 1/q_T$ in the $q_T \rightarrow 0$ limit, but this divergence can be cancelled through subtraction, using an all-orders factorisation formula for the universal q_T -dependence of the cross-section at low- q_T , which was derived to allow its resummation [116–118] and confirmed using soft-collinear effective theory (SCET) [119]. This gives the

q_T -subtraction method, schematically

$$d\hat{\sigma}_{ab \rightarrow X}^{\text{N}^k\text{LO}} = \mathcal{H}_X^{\text{N}^k\text{LO}} \otimes \frac{d\hat{\sigma}_{ab \rightarrow X}^{\text{LO}}}{dq_T} + \left[\frac{d\hat{\sigma}_{ab \rightarrow X+j}^{\text{N}^{k-1}\text{LO}}}{dq_T} - \Sigma_X^{\text{N}^k\text{LO}}(q_T) \otimes d\hat{\sigma}_{ab \rightarrow X}^{\text{LO}} \right], \quad (2.6.17)$$

where $\Sigma_X^{\text{N}^k\text{LO}}(q_T)$ contains the terms, logarithmic in q_T , which diverge in the $q_T \rightarrow 0$ limit, and so reproduces the singular behaviour of $d\hat{\sigma}_{ab \rightarrow X+j}^{\text{N}^{k-1}\text{LO}}/dq_T$ in this limit, rendering the square brackets finite (hence ‘subtraction’). The IR-finite terms, and any proportional to $\delta(q_T)$, are absorbed into the hard coefficient function $\mathcal{H}_X^{\text{N}^k\text{LO}}$, including the finite part of the k -loop virtual diagrams.

This idea can be generalised from q_T to any IR-safe resolution variable τ for which a factorisation relation exists in the $\tau \rightarrow 0$ limit. A choice which has been adopted for NNLO calculations is N -jettiness [120],

$$\tau_N = \frac{2}{Q^2} \sum_{\text{partons } k} \min_{i \in \{a, b, 1, \dots, N\}} \{q_i \cdot p_k\}, \quad (2.6.18)$$

which measures how ‘ N -jet-like’ an event is, with the q_i lightlike signal jet momenta found using a jet algorithm. In the $\tau \rightarrow 0$ limit, the event contains exactly N jets, as all parton momenta p_k are perfectly soft or aligned with a jet axis.

Such calculations are rendered challenging because the cancellation of singularities in eq. (2.6.17) is nonlocal; all singularities in phase-space are projected onto the resolution variable τ , and in practice the $\tau \rightarrow 0$ limit is typically truncated at small technical cut τ_{cut} to save computational expense and prevent numerical miscancellations of divergent quantities spoiling the result. This makes it especially important that the factors Σ_X accurately cancel the divergences of the differential cross-section, not just at the leading power but at subleading inverse powers of τ as well. The missing subleading terms are called ‘power corrections’, and have been studied in detail because of their potential to spoil the accuracy of NNLO predictions [121–123]. Their inclusion in the subtraction improves the convergence, with smaller errors for a given $\tau_{\text{cut}} > 0$.

2.6.2.2. Subtraction

The application of local subtraction to phase-space integrals is substantially more complicated than for the toy model of section 2.6.1.2, because the subtraction term must be constructed to ensure that every divergent limit of the integrand in phase-space is exactly cancelled. As seen in section 2.3.2.2, the divergent behaviour of the QCD matrix elements in these limits is highly non-trivial beyond NLO, with contributions from different colour-levels and crossings contributing (potentially-overlapping) divergences of differing degree in different kinematic regions.

Assuming such a counterterm can be constructed, the calculation can proceed, schematically, as

$$\begin{aligned}
 \int d\hat{\sigma}^{\text{NLO}} &= \int d\Phi_{n+1} d\hat{\sigma}^{\text{R}} + \int d\Phi_n d\hat{\sigma}^{\text{V}} \\
 &\equiv \int d\Phi_{n+1} \underbrace{\left[d\hat{\sigma}^{\text{R}} - d\hat{\sigma}_{\text{NLO}}^{\text{S}} \right]}_{\text{finite by universality}} \\
 &\quad + \underbrace{\int d\Phi_n \left[d\hat{\sigma}^{\text{V}} + \int d\Phi_1 d\hat{\sigma}_{\text{NLO}}^{\text{S}} \right]}_{\text{finite by KLN}}. \tag{2.6.19}
 \end{aligned}$$

A number of different methods have been developed to successfully construct these counterterms at NLO; the Catani-Seymour (CS) dipole subtraction scheme [62], and the Frixione-Kunszt-Signer (FKS) subtraction scheme [124, 125] have been most widely employed as they have been fully automated [126–130] and embedded in the popular Monte Carlo generators Sherpa [131], Herwig [132] and MadGraph [133].

At NNLO the subtler singularity structure prevents the straightforward extension of these methods to a higher order, leading to the ongoing development of alternatives. Antenna subtraction [69], which is the basis of this work, has been employed most widely and will be described in detail in chapter 3. Alternatives include sector-improved residue subtraction [134–137], nested soft-collinear subtraction [138], CoLoRFulNNLO [139] (so far applied only to electron-positron collisions), local analytic sector subtraction [140], and geometric IR subtraction [141].

CHAPTER 3

Antenna subtraction

In section 2.3.2.2 we outlined the universal behaviour of QCD colour-ordered matrix elements in their divergent limits, in which one or two real emissions could go unresolved at tree- or one-loop level.

Since this behaviour is universal, it can be extracted from physical matrix elements. This is the basis of antenna subtraction: ratios of the simplest matrix elements with the desired number and species of unresolved partons are used to define ‘antenna functions’, which are used to construct subtraction counterterms with the necessary divergent limits. It is then possible to proceed as outlined in section 2.6.2.2, constructing a subtraction counterterm to mimic the divergences associated with unresolved partons, integrating it analytically over the unresolved subspace to generate the resulting ε -poles, and adding the integrated counterterm back into a lower-multiplicity integral to cancel the poles of the loop diagrams.

The antenna functions are defined in section 3.1. The mappings that allow the final-particle phase-space to be factorised into an antenna subspace and a lower-multiplicity phase-space, required for the general analytic integration of the subtraction counterterms, are introduced in section 3.2, and the integrated antennae in section 3.3. Finally, the procedure for constructing a counterterm using these ingredients is outlined in section 3.4.

3.1. Antenna functions

As outlined above, the antenna functions are constructed from ratios of matrix elements. In each case, at least two partons must remain hard (resolved), called ‘radiators’. These radiators define the relevant reduced matrix element that will be emerge from the factorisation relation in the unresolved limit of the other particles. The expected factor accompanying it in the limit can then be extracted from the ratio. In this way they can be seen as the natural generalisation of splitting functions to general unresolved emissions.

At tree-level, this proceeds straightforwardly:

$$X_3^0(i, j, k) = S_{ijk}^{IK} \frac{M_{n+1}^{(0)}(i, j, k)}{M_n^{(0)}(I, K)} \quad (3.1.1)$$

$$X_4^0(i, j, k, l) = S_{ijkl}^{IL} \frac{M_{n+2}^{(0)}(i, j, k, l)}{M_n^{(0)}(I, L)}, \quad (3.1.2)$$

where the $S_{\dots}$ are symmetry factors accounting for identical particles and preventing double-counting. Because of the relation eq. (2.3.48), at loop-level the pure one-loop antenna is extracted by removing the factorisation of the tree-level contribution:

$$X_3^{(1)}(i, j, k) = S_{ijk}^{IK} \frac{M_{n+1}^{(1)}(i, j, k)}{M_n^{(0)}(I, K)} - X_3^0(i, j, k) \frac{M_n^{(1)}(i, j, k)}{M_n^{(0)}(I, K)} \quad (3.1.3)$$

The antenna functions can therefore be classified according to the partonic identity of the radiators: quark-antiquark $q\bar{q}$, quark-gluon qg , and gluon-gluon gg , and the radiated unresolved particles. This is the basis for the nomenclature used in tables 3.1 and 3.2 showing the set of antenna functions that can feature in NNLO QCD calculations. These were derived, respectively, from the QCD corrections to $\gamma^* \rightarrow q\bar{q}$ [142], $\tilde{\chi} \rightarrow \tilde{g}g$ [143], and $H \rightarrow gg$ [144].

As a concrete example, the tree-level three-parton quark-antiquark antenna is

$$A_3^{(0)}(1_q, 2_g, 3_{\bar{q}}) = \frac{1}{s_{123}} \left(\frac{s_{12}}{s_{23}} + \frac{s_{23}}{s_{12}} \right) + 2 \frac{s_{13}}{s_{12}s_{23}} + \mathcal{O}(\varepsilon). \quad (3.1.4)$$

Written in this form, it is easy to verify that $A_3^{(0)}$ coincides with the soft-gluon eikonal

Radiators	Radiation	Antenna functions		Ref.
		tree-level	one-loop	
$q\bar{q}$	$qg\bar{q}$	A_3^0	$A_3^1, \tilde{A}_3^1, \hat{A}_3^1$	[142]
qg	qgg $qQ\bar{Q}$	D_3^0 E_3^0	$D_3^1, \hat{D}_3^1$ $E_3^1, \tilde{E}_3^1, \hat{E}_3^1$	[143]
gg	ggg $gq\bar{q}$	F_3^0 G_3^0	$F_3^1, \hat{F}_3^1$ $G_3^1, \tilde{G}_3^1, \hat{G}_3^1$	[144]

Table 3.1: Three-parton antenna functions X_3^0 and X_3^1 for single emissions at tree- and one-loop level respectively. Tildes denote sub-leading colour contributions, suppressed by the corresponding number of inverse powers of N_c^2 , whilst hats correspond to those accompanying factors of N_F arising from a closed quark loop.

factor eq. (2.3.40) in the soft-gluon limit,

$$A_3^{(0)}(1_q, 2_g, 3_{\bar{q}}) \xrightarrow{p_2 \rightarrow 0} 2 \frac{s_{13}}{s_{12}s_{23}} \quad (3.1.5)$$

and with the collinear quark-gluon splitting function $P_g^{q(0)}(z)$ of eq. (2.1.8) in each of the collinear limits,

$$A_3^{(0)}(1_q, 2_g, 3_{\bar{q}}) \xrightarrow{p_2 \rightarrow \frac{z}{1-z} p_1} \frac{1}{s_{12}} \left[z + 2 \left(\frac{1-z}{z} \right) \right] + \mathcal{O}(s_{12}^0) \quad (3.1.6)$$

$$= \frac{1}{s_{12}} \left[\frac{1 + (1-z)^2}{z} \right]. \quad (3.1.7)$$

To disentangle the two collinear singularities, $A_3^{(0)}$ can be partitioned

$$A_3^{(0)}(1_q, 2_g, 3_{\bar{q}}) = a_3^{(0)}(1, 2, 3) + a_3^{(0)}(3, 2, 1), \quad (3.1.8)$$

in which the subantennae $a_3^{(0)}$ are constructed by partial fractioning and symmetry to contain no terms divergent in the invariant s_{jk} between the second pair of particles, and all of the terms divergent in the invariant s_{ij} between the first pair:

$$a_3^{(0)}(i, j, k) = \frac{1}{s_{ijk}} \left(\frac{s_{jk}}{s_{ij}} \right) + 2 \frac{s_{ik}}{s_{ij}(s_{ij} + s_{jk})}. \quad (3.1.9)$$

At one-loop, the antenna functions become more complicated, for instance:

$$\begin{aligned}
A_3^{(1)}(1_q, 2_g, 3_{\bar{q}}) = & 2 \left(\mathbf{I}_{qg}^{(1)}(\varepsilon, s_{12}) + \mathbf{I}_{qg}^{(1)}(\varepsilon, s_{23}) - \mathbf{I}_{q\bar{q}}^{(1)}(\varepsilon, s_{123}) \right) A_3^0(1, 2, 3) \\
& - \left(R\left(\frac{s_{12}}{s_{123}}, \frac{s_{23}}{s_{123}}\right) + \frac{5}{3} \log \frac{s_{12}}{s_{123}} + \frac{5}{3} \log \frac{s_{23}}{s_{123}} \right) A_3^0(1, 2, 3) \\
& + \frac{1}{s_{123}} + \frac{s_{13} + s_{23}}{2s_{123}s_{12}} + \frac{s_{13} + s_{12}}{2s_{123}s_{23}} \\
& - \frac{s_{12}}{2s_{123}(s_{12} + s_{13})} - \frac{s_{23}}{2s_{123}(s_{13} + s_{23})} \\
& + \frac{\log \frac{s_{12}}{s_{123}}}{s_{123}} \left(2 - \frac{1}{2} \frac{s_{12}s_{23}}{(s_{13} + s_{23})^2} + 2 \frac{s_{12} - s_{23}}{s_{13} + s_{23}} \right) \\
& + \frac{\log \frac{s_{23}}{s_{123}}}{s_{123}} \left(2 - \frac{1}{2} \frac{s_{12}s_{23}}{(s_{13} + s_{12})^2} + 2 \frac{s_{23} - s_{12}}{s_{12} + s_{13}} \right)
\end{aligned} \tag{3.1.10}$$

where the $\mathbf{I}^{(1)}$ are the Catani pole operators from eq. (2.3.28) which by eq. (2.3.27) we expect to accompany a tree-level contribution, and where

$$\begin{aligned}
R(x, y) = & \log x \log y - \log x \log(1 - x) - \log y \log(1 - y) \\
& + \frac{\pi^2}{6} - \text{Li}_2(x) - \text{Li}_2(y).
\end{aligned} \tag{3.1.11}$$

is a function collecting the terms of harmonic weight 2.¹

As might be expected, the four-parton tree-level antennae are substantially more complicated than the three-parton antennae, as they are functions of more kinematic invariants and must accommodate a larger number of possible divergent limits. They are listed in table 3.2 and can be found in full in [69].

3.1.1. Crossings

Implicitly in the above it was assumed that the collinear singularities of the radiated partons arose with final-state, rather than initial-state, radiators. This is suitable for processes that produce QCD particles from QED collisions (e.g. at an electron-positron collider), but needs some adaptation for hadronic initial-states.

¹ The emergence of non-rational functions of momentum invariants is a general phenomenon of loop integrals. It is conjectured [145] that the Laurent coefficient of ε^k of an L -loop diagram contains terms of harmonic weight no higher than $2L + k$. Here $k = 0$ and $L = 1$, and the maximal weight is indeed 2.

Radiators	Radiation	Antenna functions tree-level	Ref.
$q\bar{q}$	$qgg\bar{q}$ $q Q\bar{Q}\bar{q}$ $qq\bar{q}\bar{q}$	$A_4^0, \tilde{A}_4^0$ B_4^0 C_4^0	[142]
qg	$qggg$ $q Q\bar{Q}g$	D_4^0 $E_4^0, \tilde{E}_4^0$	[143]
gg	$gggg$ $g q\bar{q}g$ $q\bar{q} Q\bar{Q}$	F_4^0 $G_4^0, \tilde{G}_4^0$ H_4^0	[144]

Table 3.2: Four-parton antenna functions X_4^0 for single emissions at tree- and one-loop level.

The antenna functions themselves are straightforwardly obtained through crossing symmetry. Subantennae, however, are constructed to isolate only the divergences that occur in the crossed process. For example, the D_3^0 antenna is given by

$$D_3^0(1_q, 2_g, 3_g) = \frac{1}{s_{23}} \left(\frac{s_{12}}{s_{13}} + \frac{s_{13}}{s_{12}} \right) + \frac{5}{s_{123}} + \frac{s_{23}}{s_{123}^2} \left(\frac{s_{12} + s_{23}}{s_{13}} + \frac{s_{13} + s_{23}}{s_{12}} + 2 \frac{s_{12}s_{13}}{s_{23}^2} + 1 \right) + \mathcal{O}(\varepsilon) \quad (3.1.12)$$

and for configurations in which both the radiators are final-state particles, this is partitioned into two subantennae related by transposition:

$$D_3^{(0)}(1, 3, 2) = d_3^{(0)}(1, 3, 2) + d_3^{(0)}(1, 2, 3) \quad (3.1.13)$$

where

$$d_3^0(1_q, 2_g, 3_g) = \frac{s_{12}}{s_{13}s_{23}} + \frac{5}{2} \frac{1}{s_{123}} + \frac{s_{23}}{s_{123}^2} \left(\frac{s_{12} + s_{23}}{s_{13}} + \frac{s_{12}s_{13}}{s_{23}^2} + \frac{1}{2} \right) + \mathcal{O}(\varepsilon). \quad (3.1.14)$$

This is necessary because the original antenna contains both the contributions in which gluon 2_g can be identified as the radiator, and those in which 3_g is identified as the radiator. When one of the gluons has been crossed into the initial state, it can only be the radiator, as the kinematics of an incoming parton are fixed.

Instead D_3^0 must be partitioned according to the quantum numbers of the parton constructed in the collinear limit (e.g. a quark for a quark-gluon combination, or a gluon for gluon-gluon), to unambiguously identify the reduced matrix element that should accompany it. It is therefore instead decomposed as

$$D_3^0 = D_{g,gg}^0 + D_{g,gq}^0 \quad (3.1.15)$$

such that $D_{g,ij}^0$ contains only the singularities arising from the collinearity of g with i and not with j .

3.2. Factorisation of phase-space

The antenna functions defined above allow us to construct a counterterm which ensures the cancellation of the singularities in the divergent limits of phase-space. In order to proceed with the subtraction programme as outlined in eq. (2.6.19), we must be able to integrate the subtraction terms over the unresolved phase-space, mapping the $(n + m)$ particle momenta onto the n -particle resolved phase-space (which will be integrated over numerically), and integrating the antenna functions analytically over the unresolved m particles' momenta.

This requires a mapping of phase-space from the unfactorised form

$$\begin{aligned} d\Phi_{n+m}(p_1, \dots, p_i, p_{i+1}, \dots, p_{i+m+1}, \dots, p_{n+m}; Q) \\ = (2\pi)^d \delta^{(d)} \left(Q - \sum_{j=1}^{n+m} p_j \right) \prod_{j=1}^{n+m} \left(\frac{d^{d-1} p_j}{2E_j (2\pi)^{d-1}} \right) \end{aligned} \quad (3.2.1)$$

to a factorised form in which the unresolved-particle subspace is explicit,

$$\begin{aligned} d\Phi_n(p_1, \dots, p_{i-1}, q_i, q_{i+m+1}, p_{i+m+2}, \dots, p_{n+m}; P_1 + P_2) \\ \cdot d\Phi_{X_{i \dots (i+m+1)}}(p_i, \dots, p_{i+m+1}; q_i + q_{i+m+1}) \end{aligned} \quad (3.2.2)$$

for final-final antennae,

$$d\Phi_n(p_1, \dots, p_{i-1}, q_i, p_{i+m+1}, \dots, p_{n+m}; P_1 + \hat{x}_2 P_2)$$

$$\cdot \frac{1}{2\pi} \left(\sum_{a=0}^m p_{i+a} - P_2 \right)^2 d\Phi_{m+1} \left(p_i, \dots, p_{i+m}; \sum_{a=0}^m p_{i+a} \right) \delta(x_2 - \hat{x}_2) \frac{dx_2}{x_2} \quad (3.2.3)$$

for initial-final antennae, and

$$d\Phi_n(q_1, \dots, q_{i-1}, q_{i+m}, \dots, q_{n+m}; \hat{x}_1 P_1 + \hat{x}_2 P_2) \cdot \left[\prod_{a=0}^{m-1} \left(\frac{d^{d-1} p_{i+a}}{2E_{i+a} (2\pi)^{d-1}} \right) \right] \delta(x_1 - \hat{x}_1) \delta(x_2 - \hat{x}_2) dx_1 dx_2 \quad (3.2.4)$$

for initial-initial antennae. Here $\hat{x}_i$ is the unique solution to the mass-shell relation, so the delta-functions ensure that the mapped momenta are on-shell. This will be sketched below, and the details can be found in [146].

The two-particle phase-space integrates to a constant,

$$\int d\Phi_2(p_1, p_2; P_1, P_2) = \frac{1}{(4\pi)^{1-\varepsilon}} \frac{\Gamma(1-\varepsilon)}{2\Gamma(2-2\varepsilon)} (P_1 + P_2)^{-\varepsilon}, \quad (3.2.5)$$

so for $n = 2$ eq. (3.2.2) gives

$$d\Phi_{2+m}(p_1, \dots, p_{m+2}; P_1, P_2) = d\Phi_2(q_1, q_{m+2}; P_1, P_2) \cdot d\Phi_{X_{i\dots(i+m)}}(p_1, \dots, p_{m+2}; q_1 + q_{m+2}), \quad (3.2.6)$$

and so the antenna phase-space $d\Phi_{X_{i\dots(i+m)}}$ is proportional to the $(m+2)$ -particle phase-space.

Mappings $\{p_i\} \mapsto \{q_i\}$ that realise this factorisation were investigated in general in [147] and applied to antenna subtraction in [69, 146], where they can be found in full.

As an example, we illustrate the derivation of the initial-final mapping. To ensure that the final-state unresolved particles' phase-space is independent of the mapped initial-state momenta, and so can be integrated over separately, we choose a longitudinal scaling of the initial-state momentum by momentum fraction x_2 ,

$$P_2 \mapsto Q_2(x_2) = x_2 P_2. \quad (3.2.7)$$

We must then compensate by changing the momentum of the final-state radiator to

conserve momentum, which fixes its image under the mapping,

$$p_i \mapsto q_i = p_i + \sum_{a=1}^m p_{i+a} + (x_2 - 1)P_2. \quad (3.2.8)$$

We must also ensure that the mapped momentum q_i is on-shell and so can represent the composite particle in the reduced matrix elements. Imposing this gives

$$0 = q_i^2 = 2 \left[\sum_{a,b=0}^m s_{(i+a)(i+b)} + (x_2 - 1) \sum_{a=0}^m (s_{(i+a)2}) \right] \quad (3.2.9)$$

which can be solved for x_2 ,

$$x_2 = 1 - \frac{\sum_{a,b} s_{(i+a)(i+b)}}{\sum_a s_{(i+a)2}}, \quad (3.2.10)$$

to give

$$p_i \mapsto q_i = p_i + \sum_{a=1}^m p_{i+a} - \frac{\sum_{a,b} s_{(i+a)(i+b)}}{\sum_a s_{(i+a)2}} P_2. \quad (3.2.11)$$

3.3. Integrated antenna functions

With the factorisation of phase-space arising from these momentum mappings, the antenna functions can be integrated over the subspace on which they are defined. This was performed in [69, 146, 148–151], defining the integrated antenna functions

$$\mathcal{X}_{ij\dots k}^{(l)} = \left(\frac{8\pi^2}{C_\varepsilon} \right)^m \int d\Phi_{X_{ij\dots k}} X_{ij\dots k}^{(l)} \quad (3.3.1)$$

$$\mathcal{X}_{i,j\dots k}^{(l)}(x_i) = \left(\frac{8\pi^2}{C_\varepsilon} \right)^m \int d\Phi_{m+1} \frac{Q^2}{2\pi} \delta(x_i - \hat{x}_i) X_{i,j\dots k}^{(l)} \quad (3.3.2)$$

$$\mathcal{X}_{ij,k_1\dots k_m}^{(l)}(x_i, x_j) = \left(\frac{8\pi^2}{C_\varepsilon} \right)^m \int \left[\prod_{a=1}^m \left(\frac{d^{d-1}k_a}{2E_a(2\pi)^{d-1}} \right) \right] \quad (3.3.3)$$

$$x_i x_j \delta(x_i - \hat{x}_i) \delta(x_j - \hat{x}_j) X_{ij,k_1\dots k_m}^{(l)} \quad (3.3.4)$$

for final-final, initial-final, and initial-initial hard radiators respectively.

The integrated antennae contain the explicit ε -poles necessary to cancel those of the virtual loop diagrams, in combination with the mass-factorisation term. For example, applying these definitions to the three-parton tree-level antenna A_3^0 gives

[69, 146]

$$\mathcal{A}_3^0 = (Q^2)^{-\varepsilon} \left[-2\mathbf{I}_{q\bar{q}}^{(1)}(\varepsilon, Q^2) + \frac{19}{4} + \mathcal{O}(\varepsilon^1) \right] \quad (3.3.5)$$

$$= (Q^2)^{-\varepsilon} \left[\frac{1}{\varepsilon^2} + \frac{3}{2\varepsilon} + \frac{19}{4} - \frac{7\pi^2}{12} + \mathcal{O}(\varepsilon^1) \right]$$

$$\begin{aligned} \mathcal{A}_{q,gq}^0(x_i) &= -2\mathbf{I}_{q\bar{q}}^{(1)}(\varepsilon, Q^2) \delta(1-x_i) \\ &+ (Q^2)^{-\varepsilon} \left[-\frac{1}{2\varepsilon} \frac{1}{C_F} P_q^{q(0)}(x_i) + \mathcal{O}(\varepsilon^0) \right] \end{aligned} \quad (3.3.6)$$

$$\begin{aligned} \mathcal{A}_{q\bar{q},g}^0(x_i, x_j) &= -\mathbf{I}_{q\bar{q}}^{(1)}(Q^2) \delta(1-x_i) \delta(1-x_j) \\ &+ (Q^2)^{-\varepsilon} \left[-\frac{1}{2\varepsilon} \frac{1}{C_F} \left(P_q^{q(0)}(x_i) \delta(1-x_j) \right. \right. \\ &\quad \left. \left. + \delta(1-x_i) P_q^{q(0)}(x_j) \right) + \mathcal{O}(\varepsilon^0) \right] \end{aligned} \quad (3.3.7)$$

$$\mathcal{A}_{q\bar{q},q}^0(x_i, x_j) = (Q^2)^{-\varepsilon} \left[-\frac{1}{\varepsilon} P_q^{g(0)}(x_j) \delta(1-x_i) + \mathcal{O}(\varepsilon^0) \right]. \quad (3.3.8)$$

3.4. Constructing a subtraction counterterm

The ingredients of an antenna subtraction counterterm, realising in practice the subtraction concept introduced in eq. (2.6.19), have now all been introduced, and it only remains to assemble them.

3.4.1. NLO subtraction counterterms

At NLO at most one particle can become unresolved, in the tree-level diagrams of the real-emission which contribute to $d\hat{\sigma}_{ab}^R$. The infrared singularities in the unresolved limits can therefore be mimicked using only the $X_{ijk}^{(0)}$ antenna functions. They are assembled into a counterterm $d\hat{\sigma}_{ab}^{S, \text{NLO}}$ for each colour-ordered squared-matrix-element $M_{n+1}^{(0)}$. At the leading-order of the colour expansion,

$$d\hat{\sigma}_{ab}^R \propto \sum_{\sigma} d\Phi_{n+1} M_{n+1}^{(0)}(\sigma(\{p_i\})) J_n^{n+1}(\sigma(\{p_i\})) \quad (3.4.1)$$

and so the counterterm is constructed as

$$d\hat{\sigma}_{ab}^{\text{SFF NLO}} \propto \sum_{\sigma} d\Phi_{n+1} \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} X_{(j-1)j(j+1)}^0 \quad (3.4.2)$$

$$\times M_n^0(p_1, \dots, p_{j-2}, q_{j-1}, q_{j+1}, p_{j+2}, \dots, p_{n+1}) J_n^n$$

$$d\hat{\sigma}_{\hat{a}\hat{b}}^{\text{SIF NLO}} \propto \sum_{\sigma} d\Phi_{n+1} \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} X_{\hat{a},j(j+1)}^0 \quad (3.4.3)$$

$$\times M_n^0(\dots, p_{j-2}, q_{j-1}, p_{j+1}, p_{j+2}, \dots; x_{\hat{a}} P_{\hat{a}}) J_n^n$$

$$d\hat{\sigma}_{\hat{a}\hat{b}}^{\text{SII NLO}} \propto \sum_{\sigma} d\Phi_{n+1} \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} X_{\hat{a}\hat{b},j}^0 \quad (3.4.4)$$

$$\times M_n^0(q_1, \dots, q_{j-1}, q_{j+1}, \dots, q_n; x_{\hat{a}} P_{\hat{a}}, x_{\hat{b}} P_{\hat{b}}) J_n^n$$

for final-final, initial-final, and initial-initial radiators respectively, where the hat denotes the initial-state radiator. The jet functions depend only on the resolved and mapped momenta. Where one or both of the colliding particles are leptons rather than hadrons (such as for e^+e^- or DIS), the corresponding initial-state-radiator subtraction terms are not necessary, and are zero. These subtraction terms integrate over the factorised antenna phase-space to give, respectively (with an additional, conventional, minus-sign),

$$d\hat{\sigma}_{ab}^{\text{TFF NLO}} \propto - \sum_{\sigma} d\Phi_n \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} \mathcal{X}_{(j-1)j(j+1)}^0 M_n^0 J_n^n \quad (3.4.5)$$

$$d\hat{\sigma}_{\hat{a}\hat{b}}^{\text{TIF NLO}} \propto - \sum_{\sigma} d\Phi_n \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} \int \frac{dx_{\hat{a}}}{x_{\hat{a}}} \mathcal{X}_{\hat{a},j(j+1)}^0(x_{\hat{a}}) M_n^0 J_n^n \quad (3.4.6)$$

$$d\hat{\sigma}_{\hat{a}\hat{b}}^{\text{TII NLO}} \propto - \sum_{\sigma} d\Phi_n \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} \int \frac{dx_{\hat{a}}}{x_{\hat{a}}} \frac{dx_{\hat{b}}}{x_{\hat{b}}} \mathcal{X}_{\hat{a}\hat{b},j}^0(x_{\hat{a}}, x_{\hat{b}}) M_n^0 J_n^n. \quad (3.4.7)$$

The ε -poles of the virtual contribution then cancel directly against those of the integrated subtraction counterterm and the mass-factorisation contribution of eq. (2.3.56), which can be verified explicitly. The partonic cross-section $d\hat{\sigma}_{ab}^{\text{NLO}}$ can

then be computed as

$$\begin{aligned}
\int d\hat{\sigma}_{ab}^{\text{NLO}} &= \int_{n+1} d\hat{\sigma}_{ab}^{\text{R}} + \int_n d\hat{\sigma}_{ab}^{\text{V}} \\
&\equiv \int_{n+1} \left[d\hat{\sigma}_{ab}^{\text{R}} - \sum_{\text{II,IF,FF}} d\hat{\sigma}_{ab}^{\text{S NLO}} \right] \\
&\quad + \int_n \left[d\hat{\sigma}_{ab}^{\text{V}} - \sum_{\text{II,IF,FF}} d\hat{\sigma}_{ab}^{\text{T NLO}} + d\hat{\sigma}_{ab}^{\text{MF NLO}} \right]. \tag{3.4.8}
\end{aligned}$$

3.4.2. NNLO subtraction counterterms

The counterterm structure at NNLO is more complicated and we will only summarise its formation. The ideas behind its structure are similar to those at NLO, but are naturally complicated by the presence of double-unresolved limits in the double-real contribution, and the emergence of both real- and virtual-type IR singularities in the real-virtual. Clearly, the NNLO subtraction must contain the terms that would provide the NLO subtraction for the final-state with an extra required jet. A comprehensive explanation is given in [64], and will only be summarised here. As for the toy-model, however, the subtraction terms needed to reproduce the correct limiting behaviour but are otherwise not unique: the systematic approach of following colour-orderings is simply one way to achieve this.

Overall, at NNLO the partonic cross-section can be expressed as

$$\begin{aligned}
\int d\hat{\sigma}_{ab}^{\text{NNLO}} &= \int_{n+2} d\hat{\sigma}_{ab}^{\text{RR}} + \int_{n+1} d\hat{\sigma}_{ab}^{\text{RV}} + \int_n d\hat{\sigma}_{ab}^{\text{VV}} \\
&\equiv \int_{n+2} \left[d\hat{\sigma}_{ab}^{\text{RR}} - \sum_{\text{II,IF,FF}} d\hat{\sigma}_{ab}^{\text{S}} \right] \\
&\quad + \int_{n+1} \left[d\hat{\sigma}_{ab}^{\text{RV}} - \sum_{\text{II,IF,FF}} d\hat{\sigma}_{ab}^{\text{T}} + d\hat{\sigma}_{ab}^{\text{MF RV}} \right] \\
&\quad + \int_n \left[d\hat{\sigma}_{ab}^{\text{VV}} - \sum_{\text{II,IF,FF}} d\hat{\sigma}_{ab}^{\text{U}} + d\hat{\sigma}_{ab}^{\text{MF VV}} \right]. \tag{3.4.9}
\end{aligned}$$

3.4.2.1. Double-real subtraction

The double-real subtraction term for an $X + (n \text{ jet})$ final-state must subtract divergences arising from:

- (a) a single unresolved parton (with n identified jets);
- (b) two unresolved partons, consecutive in the colour ordering, between the same pair of hard radiators (‘colour-connected’);
- (c) two unresolved partons sharing a common radiator (‘almost colour-connected’);
- (d) two unresolved partons with distinct radiators (‘colour-disconnected’).

Subtraction terms for (a) are built exactly as the NLO real-radiation subtraction for the $(n + 1)$ -jet final-state would be (with the only difference the number of jets allowed by the jet-function, J_n^{n+1} rather than J_{n+1}^{n+1}). Those for (b) are built from the four-parton antenna functions in the same way (combined with products of two three-parton antennae to prevent the oversubtraction of single-unresolved limits in the antenna functions themselves).

These are sufficient for hadronic-collision processes with colourless final-states X at leading-order, because the double-real contribution then contains four coloured partons, of which only the two final-state real emissions can become unresolved. They therefore fall into categories (a) or (b) above. For processes with a single final-state parton at leading order, (c) can arise, and for those with two (d) can arise. In each case additional terms must be built, from products of three-parton antenna functions, to compensate for over-subtraction of parts (a) and (b), disentangling the divergent limits where a parton is resolved with respect to one radiator but not another. The full details can be found in [64].

3.4.2.2. Real-virtual subtraction

The real-virtual subtraction term must subtract divergences arising from:

- (a) the explicit ε -poles of the real-virtual matrix elements;

(b) the single-unresolved limits of the real-virtual matrix elements.

Exactly as at NLO, the (a) contribution to the double-real radiation subtraction term integrates and combines with the $(n+3)$ -parton matrix elements of the real-virtual mass-factorisation contribution to exactly cancel the ε -poles arising from the loop matrix elements, (a) above. This cancellation must occur exactly as at NLO, as it is simply the NLO calculation for the $X + \text{jet}$ calculation (as only the jet function differs, as above, which doesn't affect the ε -poles).

To build a subtraction term for (b), we use the definition of the antenna in eq. (3.1.3) and the factorisation behaviour of eq. (2.3.48) to construct the subtraction term. For example, for final-final radiators,

$$\begin{aligned} d\hat{\sigma}_{ab}^{\text{T}_{\text{FF}}} \propto \sum_{\sigma} d\Phi_{n+1} \sum_{\substack{\text{unres.} \\ \text{parton} \\ j}} \left[X_{(j-1)j(j+1)}^0 M_n^{(1)} \right. \\ \left. + \left(X_{(j-1)j(j+1)}^1 - \beta_0 X_{(j-1)j(j+1)}^0 \log \frac{|s_{(j-1)j(j+1)}|}{\mu_{\text{R}}^2} \right) M_n^{(0)} \right] J_n^n, \end{aligned} \quad (3.4.10)$$

where the term proportional to β_0 is required to correct for the fact that the loop of the antenna function X_{ijk}^1 is renormalised at scale s_{ijk} , whilst the one-loop matrix element is renormalised at scale μ_{R} .

For final-states containing one or more coloured partons at leading-order, further terms are once more required, to compensate for the (c) and (d) contributions of the double-real subtraction term.

3.4.2.3. Double-virtual subtraction

The double-virtual two-loop matrix elements occupy the Born phase space, and so there is no opportunity for final-state partons to become unresolved. As a result, the double-virtual subtraction term is entirely constrained by the book-keeping of the previous contributions, as it must balance any terms introduced above that do not already cancel between the double-real and real-virtual subtractions. The resulting structure can be found in [64].

Channel order	Partonic channel		
	$q_f \bar{q}_f, \bar{q}_f q_f$	$qg, gq, \bar{q}g, g\bar{q}$	$gg, qq, q\bar{q}, \bar{q}q, \bar{q}\bar{q}$
LO	α_s^0	α_s^1	α_s^2
NLO	α_s^1	α_s^2	α_s^3
NNLO	α_s^2	α_s^3	α_s^4

Table 3.3: The order in the strong coupling at which each partonic channel contributes to the $pp \rightarrow \gamma\gamma$ cross-section. As a result, for a NNLO calculation of the latter, at $\mathcal{O}(\alpha_s^2)$, only the (same-flavour) $q\bar{q}$ and $\bar{q}q$ channels must themselves be calculated to NNLO.

The calculation can then be checked with the requirement that the ε -poles of the integrated subtraction terms exactly cancel the ε -poles of the two-loop matrix elements, by reproducing the Catani pole operators eq. (2.3.33) up to $\mathcal{O}(\varepsilon^0)$.

3.5. Application to diphoton production

The quark lies in the intersection of particles charged under QCD and QED. As a result, the respective gauge bosons do not couple directly (i.e. there is no gluon-photon vertex) but through quarks as an intermediary. As a result, counting the perturbative order of contributions to the cross-section consistently across the partonic contributions, at LO only the $q\bar{q}$ -channel contributes, whilst at NLO a single initial-state gluon is possible, and at NNLO two initial-state gluons become possible. The relative orders of contributions in each channel are shown explicitly in table 3.3. Illustrative Feynman diagrams for each of these channels are shown in fig. 3.1.

For concreteness, we will summarise the calculation, including the calculation of the amplitudes and matrix elements, and the application of the antenna formalism.

3.5.1. LO

At leading-order, the only Feynman diagram that contributes is the first diagram of fig. 3.1. Since there are no gluons, this is effectively a pure-QED diagram for each quark colour, there is trivial colour structure, and it carries strong-coupling factor α_s^0 (i.e. $m = 0$ in eq. (2.2.1)). The diagram can easily be computed with Feynman diagram techniques to give

$$B_{2;2\gamma}^0(q, \bar{q}; \gamma_1, \gamma_2) = 4 \frac{\sum_{i=1}^2 [s_{q\gamma_i}^3 s_{\bar{q}\gamma_i} + s_{q\gamma_i} s_{\bar{q}\gamma_i}^3]}{\prod_{i=1}^2 s_{q\gamma_i} s_{\bar{q}\gamma_i}} \quad (3.5.1)$$

$$= 8 \left[\frac{s_{q\gamma_1}}{s_{q\gamma_2}} + \frac{s_{q\gamma_2}}{s_{q\gamma_1}} \right], \quad (3.5.2)$$

where we adopt the notation $B_{m;2\gamma}^l$ for the leading-colour colour-ordered matrix element $M_{m+2\gamma}^l$ with l loops, 2 photons, and m QCD partons of which two are a quark-antiquark pair. This matrix element must then be dressed with gauge factors and averaged over initial states (of definite quark flavour, for convolution with the appropriate PDFs) to give a factor we will denote $N_{q\bar{q}}^{\text{LO}}$ for later reference,

$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{LO}} = N_{q\bar{q}}^{\text{LO}} \delta_{f_2}^{f_1} \left[d\Phi_2 B_{2;2\gamma}^0 \right] J_0^0 \quad (3.5.3)$$

$$:= \underbrace{\left(\frac{1}{2N_c} \right)^2}_{\text{average over initial states}} \underbrace{\frac{1}{2!}}_{\text{identical FS particles}} \underbrace{(4\pi\alpha_{\text{em}})^2}_{\text{QED coupling (for } 2\gamma)} \underbrace{\sum_{i,j=1}^{N_c} \delta_{ij}}_{N_c} \underbrace{Q_f^4}_{\text{EM quark charge}} \underbrace{\delta_{f_2}^{f_1}}_{\text{matching quark flavours}} \left[d\Phi_2 B_{2;2\gamma}^0 \right]. \quad (3.5.4)$$




Figure 3.1: Example Born, real and double-real diagrams for diphoton production, illustrating the opening of new partonic channels at higher orders. Various other diagrams arise as crossings of these.

Here f_1 and f_2 indicate the quark flavours for convolution with the appropriate PDFs. The matrix element $B_{2;2\gamma}^0$ is only non-zero if the two external quarks are of the same quark flavour, which is imposed by the $\delta_{f_2}^{f_1}$ factor, and so combining eqs. (2.1.3) and (2.2.1),

$$\begin{aligned} \sigma_{pp \rightarrow \gamma\gamma}^{\text{LO}} = N_{q\bar{q}}^{\text{LO}} \sum_f Q_f^4 \int dx_1 dx_2 d\Phi_2 f_{q_f}^p(x_1; \mu_F^2) f_{\bar{q}_f}^p(x_2; \mu_F^2) B_{2;2\gamma}^0(1_q, 2_{\bar{q}}; 3_\gamma, 4_\gamma) \\ + \mathcal{O}(\alpha_s) \end{aligned} \quad (3.5.5)$$

is the leading-order contribution to the cross-section for diphoton production.

Although the matrix element $B_{2;2\gamma}^0$ is a tree-level matrix element, and therefore finite without requiring dimensional regularisation, it will appear within subsequent virtual corrections according to the Catani operators eqs. (2.3.27) and (2.3.33), calculated in dimensional regularisation.² The $\mathcal{O}(\varepsilon^l)$ terms therefore become relevant to the finite part of an l -loop matrix element. We therefore quote the matrix element to $\mathcal{O}(\varepsilon)$ for later reference from [155],

$$B_{2;2\gamma}^0 = 8 \left[\frac{s_{q\gamma_1}}{s_{q\gamma_2}} + \frac{s_{q\gamma_2}}{s_{q\gamma_1}} \right] - 16\varepsilon \left[\frac{s_{q\gamma_1}}{s_{q\gamma_2}} + \frac{s_{q\gamma_2}}{s_{q\gamma_1}} + 1 \right] + \mathcal{O}(\varepsilon^2). \quad (3.5.6)$$

3.5.2. NLO

At NLO the $q\bar{q}$ channel receives real-emission corrections from the radiation of a gluon, and corresponding virtual corrections corresponding to the emission and reabsorption of a virtual gluon. As described in section 2.2, these both contribute factors of α_s to the overall cross-section, which will accompany the partonic cross-sections $d\hat{\sigma}_{ab}^{\text{NLO}}$ according to eq. (2.2.1).

² The results quoted throughout were calculated in ‘conventional dimensional regularisation’ (CDR) [152], in which all fields are taken to be D -dimensional; the tree-level contributions must therefore also be calculated in CDR for the relationships between loop- and tree-level matrix elements given by the Catani operators eqs. (2.3.27) and (2.3.33) to hold. The main alternatives to CDR are the original formulation of dimensional regularisation, the ‘t Hooft–Veltman (HV) scheme [153], which treats the external states as 4-dimensional and only the internal states as D -dimensional, and the ‘dimensional reduction’ (DRED) scheme [154], in which only the momenta of the unobserved particles are D -dimensional, and all other momenta, helicities and polarisation vectors are 4-dimensional.

Again, this diagram is computable with Feynman diagram techniques to give

$$B_{3;2\gamma}^0(q, g, \bar{q}; \gamma_1, \gamma_2) = 4s_{q\bar{q}} \left[\frac{(s_{qg}^3 s_{\bar{q}g} + s_{qg} s_{\bar{q}g}^3) + \sum_{i=1}^2 [s_{q\gamma_i}^3 s_{\bar{q}\gamma_i} + s_{q\gamma_i} s_{\bar{q}\gamma_i}^3]}{s_{qg} s_{\bar{q}g} \prod_{i=1}^2 s_{q\gamma_i} s_{\bar{q}\gamma_i}} \right] + \mathcal{O}(\varepsilon) \quad (3.5.7)$$

where the interchange symmetry between the photons is explicit, and the eikonal factor of eq. (2.3.40) can be read directly from the ratio to eq. (3.5.1) in the divergent limits. This must again be dressed with colour and symmetry factors, to give

$$\begin{aligned} d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^R &= 8\pi^2 N_{q\bar{q}}^{\text{LO}} Q_f^4 \frac{N_c^2 - 1}{N_c} \delta_{f_2}^{f_1} \left[d\Phi_3 B_{3;2\gamma}^0 \right] J_0^1 \\ &= 8\pi^2 \left(\frac{1}{2N_c} \right)^2 \frac{1}{2!} (4\pi\alpha_{\text{em}})^2 Q_f^4 (N_c^2 - 1) \delta_{f_2}^{f_1} \left[d\Phi_3 B_{3;2\gamma}^0 \right]. \end{aligned} \quad (3.5.8)$$

The factor of $8\pi^2$ converts between the expansion in powers of

$$g_s = \sqrt{4\pi\alpha_s} \quad (3.5.9)$$

which arise directly from the Feynman rules, and the convention adopted for the perturbative expansion of eq. (2.2.1) in powers of $\alpha_s/2\pi$; the matrix elements derived from the Feynman rules can then be used directly.

We can now see how antenna subtraction works in practice. Since the only QCD partons in the matrix element are a quark-antiquark pair and a gluon, the subtraction term is constructed with the A_3^0 antenna of eq. (3.1.4). Only the final-state gluon can become unresolved, so the two radiators must be the initial state quark and antiquark, and hence

$$d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{SFF NLO}} = 0 \quad (3.5.10)$$

$$d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{SIF NLO}} = 0 \quad (3.5.11)$$

$$\begin{aligned} d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{SII NLO}} &= 8\pi^2 N_{q\bar{q}}^{\text{LO}} Q_f^4 \frac{N_c^2 - 1}{N_c} \delta_{f_2}^{f_1} \\ &\times \left[d\Phi_3 A_3^0(\hat{1}_q, g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) \right], \end{aligned} \quad (3.5.12)$$

where the primed momenta are the images of the momentum maps of section 3.2. It

is now clear from the properties of the antenna functions outlined in section 3.1 that

$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^R - d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{S\text{ NLO}} &\propto \delta_{f_2}^{f_1} d\Phi_3(p_{\gamma_1}, p_{\gamma_2}, p_g) \left[B_{3;2\gamma}^0(1_q, g, 2_{\bar{q}}; 3_\gamma, 4_\gamma) J_0^1 \right. \\ &\quad \left. - A_3^0(\hat{1}_q, g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) J_0^0 \right] \end{aligned} \quad (3.5.13)$$

is indeed the integral of a finite quantity over a space with integer dimension and can be evaluated numerically. This was the goal of the subtraction procedure.

As discussed, the $q\bar{q}$ channel contributes, at its leading order but at NLO in the global expansion, through the second diagram of fig. 3.1. Since this is a crossing of the diagram for $q\bar{q} \rightarrow \gamma\gamma g$, the matrix element can be reused with a permutation of the arguments, dressed with the appropriate colour factors and averages. Using the same arguments as in eq. (3.5.4), with an additional factor arising from the Fierz identity eq. (2.3.38) and the sum over final spin states of the additional particle,

$$2 \sum_{i,j=1}^{N_c} T_{ij}^a T_{ji}^a = 2 \cdot \frac{1}{2} \sum_{i,j=1}^{N_c} \left[\delta_{ii} \delta_{jj} - \frac{1}{N_c} \delta_{ij} \delta_{ji} \right] = N_c^2 - 1 \quad (3.5.14)$$

the prefactor is

$$N_{q\bar{q}}^{\text{NLO}} = \left(\frac{1}{2^2 N_c (N_c^2 - 1)} \right) \frac{1}{2!} (4\pi\alpha_{\text{em}})^2 Q_q^4 (N_c^2 - 1) = N_{q\bar{q}}^{\text{LO}}, \quad (3.5.15)$$

to give

$$d\hat{\sigma}_{q\bar{q}}^R = 8\pi^2 N_{q\bar{q}}^{\text{NLO}} Q_q^4 \left[d\Phi_3 B_{3;2\gamma}^0(1'_q, 2_g, \bar{q}; 3_\gamma, 4_\gamma) \right]. \quad (3.5.16)$$

Here and throughout the quark flavour index of the $q\bar{q}$ -channel is suppressed wherever dependence upon it is trivial, with the final-state quark flavour always equal to the initial-state quark flavour.

Again the radiators must both be in the initial-state, so

$$d\hat{\sigma}_{qg}^{S_{\text{FF}}\text{ NLO}} = 0 \quad (3.5.17)$$

$$d\hat{\sigma}_{qg}^{S_{\text{IF}}\text{ NLO}} = 0 \quad (3.5.18)$$

$$d\hat{\sigma}_{qg}^{S_{\text{II}}\text{ NLO}} = 8\pi^2 N_{qg}^{\text{NLO}} Q_q^4 \left[d\Phi_3 A_3^0(\hat{1}_q, \hat{2}_g, \bar{q}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) \right], \quad (3.5.19)$$

where again it is now easy to verify that the integral to be evaluated numerically

$$d\hat{\sigma}_{gg}^R - d\hat{\sigma}_{gg}^{S\text{ NLO}} \quad (3.5.20)$$

has been rendered finite by the counterterm. The partonic gq , $\bar{q}g$ and $g\bar{q}$ cross-sections all follow by symmetry, only differing in the final convolution with their respective PDFs.

We can now turn to the virtual and integrated-subtraction term contributions, $d\hat{\sigma}_{ab}^V$ and $d\hat{\sigma}_{ab}^{T\text{ NLO}}$. Since the integrated subtraction terms are the integrals of the subtraction terms, we can write them down immediately from eqs. (3.5.10) to (3.5.12) and eqs. (3.5.17) to (3.5.19):

$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{T_{\text{FF}}\text{ NLO}} = d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{T_{\text{IF}}\text{ NLO}} = 0 \quad (3.5.21)$$

$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{T_{\text{II}}\text{ NLO}} = N_{q\bar{q}}^{\text{LO}} Q_f^4 \frac{N_c^2 - 1}{N_c} \delta_{f_2}^{f_1} \quad (3.5.22)$$

$$\times \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \mathcal{A}_{q\bar{q},g}^0(x_1, x_2) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) \right]$$

$$d\hat{\sigma}_{qg}^{T_{\text{FF}}\text{ NLO}} = d\hat{\sigma}_{qg}^{T_{\text{IF}}\text{ NLO}} = 0 \quad (3.5.23)$$

$$d\hat{\sigma}_{qg}^{T_{\text{II}}\text{ NLO}} = N_{qg}^{\text{NLO}} Q_q^4 \quad (3.5.24)$$

$$\times \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \mathcal{A}_{q\bar{g},\bar{q}}^0(x_1, x_2) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) \right],$$

where the factors of $(8\pi^2)$ have been absorbed into the definition of the integrated antenna function, as in eq. (3.3.3), and C_ε is the geometric factor for loop integrals defined in eq. (2.3.9).

Since the qg channel first contributes at tree-level at this order, the virtual contribution must be zero,

$$d\hat{\sigma}_{qg}^V = 0. \quad (3.5.25)$$

As a result, the poles introduced by the integrated antenna functions in eq. (3.5.24) can only cancel against the mass-factorisation contribution of eq. (2.3.56). Using eq. (3.3.8) we can extract the principal part of the Laurent series in ε of the integrated

subtraction term,

$$\begin{aligned} \text{P} \left[d\hat{\sigma}_{qg}^{\text{T}_{\text{H}}\text{NLO}} \right] &= \left(\frac{Q_q^4}{2^3 N_c} \right) (4\pi\alpha_{\text{em}})^2 \\ &\times \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \left[-\frac{1}{\varepsilon} P_q^{g(0)}(x_2) \delta(1-x_1) \right] B_{2;2\gamma}^0 \right]. \end{aligned} \quad (3.5.26)$$

Likewise, using the expression for the NLO mass-factorisation contribution eq. (2.3.56) and the leading-order cross-section from eq. (3.5.3), we can construct the mass-factorisation counterterm for the qg -channel,

$$\begin{aligned} d\hat{\sigma}_{qfg}^{\text{MF NLO}} &= -C_\varepsilon \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \left[\delta(1-x_1) d\hat{\sigma}_{qf\bar{q}f}^{\text{LO}} \Gamma_{\bar{q}fg}^1(x_2) \right] \\ &= N_{q\bar{q}}^{\text{LO}} Q_f^4 C_\varepsilon \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \delta(1-x_1) \Gamma_{qg}^1(x_2) B_{2;2\gamma}^0 \right] \\ &= \left(\frac{Q_f^4 C_\varepsilon}{2^3 N_c} \right) (4\pi\alpha_{\text{em}})^2 \\ &\times \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \delta(1-x_1) \left[-\frac{1}{\varepsilon} P_{\bar{q}}^{g(0)}(x_2) \right] B_{2;2\gamma}^0 \right]. \end{aligned} \quad (3.5.27)$$

Comparing the two (and noting that $P_{\bar{q}}^g(x) \equiv P_q^g(x)$), we can verify that for the qg -channel the ε -pole in the virtual subtraction term does indeed cancel against the mass-factorisation contribution, and so

$$\lim_{\varepsilon \rightarrow 0} \left[d\hat{\sigma}_{qg}^{\text{V}} - d\hat{\sigma}_{qg}^{\text{T}_{\text{H}}\text{NLO}} + d\hat{\sigma}_{qg}^{\text{MFNLO}} \right] = \text{finite}. \quad (3.5.28)$$

This was the end goal of the antenna subtraction procedure: the divergences associated with unresolved real emissions have been regulated, the corresponding integrals performed analytically in dimensional-regularisation over the antenna phase-space, and cancelled as poles in ε against those arising elsewhere in the calculation.

We now repeat this process for the $q\bar{q}$ channel. The virtual matrix elements were computed in [155]. For brevity we omit here some terms of $\mathcal{O}(\varepsilon^1)$ or higher:

$$\begin{aligned} B_{2;2\gamma}^1(q, \bar{q}; \gamma_1, \gamma_2) &= \left(-\frac{1}{\varepsilon} + 4 - 5\varepsilon \right) B_{2;2\gamma}^0 \text{Bub} \left(\frac{s_{q\bar{q}}}{\mu_{\text{R}}^2} \right) + 16(2 - 5\varepsilon) \text{Bub} \left(\frac{s_{q\bar{q}}}{\mu_{\text{R}}^2} \right) \\ &+ 4 \left[3 \frac{s_{q\bar{q}}}{s_{q\gamma_1}} + 1 + \varepsilon \left(5 + \frac{s_{q\gamma_2}}{s_{q\gamma_1}} \right) \right] \text{Bub} \left(\frac{s_{q\gamma_1}}{\mu_{\text{R}}^2} \right) \\ &+ 4 \left[3 \frac{s_{q\bar{q}}}{s_{q\gamma_2}} + 1 + \varepsilon \left(5 + \frac{s_{q\gamma_1}}{s_{q\gamma_2}} \right) \right] \text{Bub} \left(\frac{s_{q\gamma_2}}{\mu_{\text{R}}^2} \right) \end{aligned}$$

$$\begin{aligned}
& + 8 \left[\frac{s_{q\bar{q}}^2 + s_{q\gamma_1}^2}{s_{q\gamma_1}} \right] \text{Box}^6(s_{q\bar{q}}, s_{q\gamma_1}) + 8 \left[\frac{s_{q\bar{q}}^2 + s_{q\gamma_2}^2}{s_{q\gamma_2}} \right] \text{Box}^6(s_{q\bar{q}}, s_{q\gamma_2}) \\
& + \mathcal{O}(\varepsilon^1)
\end{aligned} \tag{3.5.29}$$

where the Bubble integral is that of eq. (2.3.2), in the $\overline{\text{MS}}$ scheme and with an additional conventional factor of $4\pi^2$,

$$\begin{aligned}
\text{Bub} \left(\frac{s}{\mu_R^2} \right) &= e^{\varepsilon\gamma_E} \Gamma(\varepsilon) \left(-\frac{s}{\mu_R^2} \right)^{-\varepsilon} B(1-\varepsilon, 1-\varepsilon) \\
&= \frac{e^{\varepsilon\gamma_E}}{\Gamma(1-\varepsilon)} \left[\frac{1}{\varepsilon} + 2 + \mathcal{O}(\varepsilon^2) \right] \left(-\frac{s}{\mu_R^2} \right)^{-\varepsilon}
\end{aligned} \tag{3.5.30}$$

and Box^6 is the one-loop box integral evaluated in $d = 6 - 2\varepsilon$ dimensions (and therefore finite as $\varepsilon \rightarrow 0$), which can be found in eqs. (C.1–2) of [155]. From this we can calculate, using the $\mathcal{O}(\varepsilon^1)$ terms from eq. (3.5.3),

$$\text{P} \left[B_{2;2\gamma}^1 \right] = \frac{e^{\varepsilon\gamma_E}}{\Gamma(1-\varepsilon)} \left[-\frac{1}{\varepsilon^2} B_{2;2\gamma}^0 + \frac{1}{\varepsilon} \left[\left(-\frac{3}{2} + \log \left(-\frac{s_{q\bar{q}}}{\mu_R^2} \right) \right) B_{2;2\gamma}^0 \right] \right]. \tag{3.5.31}$$

$B_{2;2\gamma}^1$ must be dressed with the appropriate colour, charge and symmetry factors to give the virtual partonic cross-section,

$$\text{d}\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{V}} = N_{q\bar{q}}^{\text{LO}} Q_f^4 \frac{N_c^2 - 1}{N_c} \delta_{f_2}^{f_1} \left[\text{d}\Phi_2 B_{2;2\gamma}^1 \right]. \tag{3.5.32}$$

The ε -poles in the integrated subtraction term of eq. (3.5.22) can be extracted from the integrated antenna function, eq. (3.3.7) and the expression for the Catani operator of eq. (2.3.32), to give

$$\begin{aligned}
\text{P} \left[\text{d}\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{TII NLO}} \right] &= N_{q\bar{q}}^{\text{LO}} Q_f^4 \frac{N_c^2 - 1}{N_c} \delta_{f_2}^{f_1} \\
&\times \left[-\frac{e^{\varepsilon\gamma_E}}{\Gamma(1-\varepsilon)} \left[\frac{1}{\varepsilon^2} + \frac{3}{2\varepsilon} \right] \left(-\frac{s_{q\bar{q}}}{\mu_R^2} \right)^{-\varepsilon} \text{d}\Phi_2 B_{2;2\gamma}^0 \right. \\
&\quad \left. + \left[-\text{d}\Phi_2 \frac{1}{C_F} \int \frac{\text{d}x}{x} \frac{1}{\varepsilon} P_q^{q(0)}(x) B_{2;2\gamma}^0 \right] \right].
\end{aligned} \tag{3.5.33}$$

Finally the mass-factorisation contribution, using eq. (2.3.56), is

$$\text{d}\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{MF NLO}} = -C_\varepsilon \int \frac{\text{d}x_1}{x_1} \frac{\text{d}x_2}{x_2} \left[\delta(1-x_1) \text{d}\hat{\sigma}_{q_{f_1}\bar{q}_{f_1}}^{\text{LO}} \Gamma_{\bar{q}_{f_1}\bar{q}_{f_2}}^1(x_2) \right]$$

$$\begin{aligned}
& + \Gamma_{q_{f_2} q_{f_1}}^1(x_1) \, d\hat{\sigma}_{q_{f_2} \bar{q}_{f_2}}^{\text{LO}} \, \delta(1-x_2) \Big] \\
& = N_{q\bar{q}}^{\text{LO}} Q_f^4 C_\varepsilon \, \delta_{f_2}^{f_1} \left[-d\Phi_2 \int \frac{dx}{x} \frac{2}{\varepsilon} P_q^{(0)}(x) B_{2;2\gamma}^0 \right]. \tag{3.5.34}
\end{aligned}$$

Combining these pieces we can see that again the poles cancel in the sum,

$$\lim_{\varepsilon \rightarrow 0} \left[d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{V}} - d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{TII NLO}} + d\hat{\sigma}_{q_{f_1} \bar{q}_{f_2}}^{\text{MF NLO}} \right] = \text{finite}. \tag{3.5.35}$$

3.5.3. NNLO

The process-specific ingredients required to calculate diphoton production at NNLO are:

- the tree-level, six-parton matrix elements that provide the double-real correction, calculated in [156] (both with a single quark pair, $B_{4;2\gamma}^0$ and $\tilde{B}_{4;2\gamma}^0$, and with two quark pairs, $C_{4;2\gamma}^0$ and $D_{4;2\gamma}^0$);
- the one-loop, five-parton real-virtual matrix elements $B_{3;2\gamma}^1$, $\tilde{B}_{3;2\gamma}^1$, $\hat{B}_{3;2\gamma}^1$ and $\bar{B}_{3;2\gamma}^1$, calculated in [157] and simplified in [158];
- the two-loop, four-parton double-virtual matrix elements for $q\bar{q} \rightarrow \gamma\gamma$, $B_{2;2\gamma}^2$, $\tilde{B}_{2;2\gamma}^2$ and $\hat{B}_{2;2\gamma}^2$, calculated in [155];
- the one-loop, four-parton ‘gluon-gluon box’ matrix element for $gg \rightarrow \gamma\gamma$ via a quark loop, $A_{2;2\gamma}^1$, calculated in [159].

These will be introduced in detail below. For concreteness we will focus on the $q\bar{q}$ -channel, which, being fully-NNLO, introduces additional subtleties to the subtraction method in the double-real contribution.

As previously discussed, obtaining a photonic final-state from a QCD initial-state requires a quark line, which constrains the possible diagrams at lower orders to have as external particles the two photons, a single quark-antiquark pair, and up to one gluon. At NNLO, there can be a second emitted gluon (giving $B_{4;2\gamma}^0$, as depicted in fig. 3.2a), or a virtual gluon can split into a second quark-antiquark pair, with the emission of the two photons possible from either the initial-state or the final-state quark line. Similarly, in a crossing of this Feynman diagram it becomes possible




Figure 3.2: Example double-real radiation diagrams, depicted in the $q\bar{q}$ channel, contributing to $d\hat{\sigma}_{q\bar{q}}^{\text{RR}}$.

for the initial-state quark-antiquark pair to propagate to the final-state, interacting only through the exchange of a t -channel gluon. Again, any combination of photonic emissions from the two quark lines is possible. Since they can represent quarks of different flavours, the electromagnetic charge factor Q_q can no longer be globally extracted from the matrix elements.

3.5.3.1. Double-real

We can write the double-real contribution to the $q\bar{q}$ -channel as

$$\begin{aligned}
 d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{RR}} &= \left(8\pi^2\right)^2 N_{q\bar{q}}^{\text{LO}} \left(N_c^2 - 1\right) d\Phi_4 \\
 &\times \left[\delta_{f_2}^{f_1} \frac{Q_f^4}{2} \left[\left(B_{4;2\gamma}^0 \left(1_q, 3_g, 4_g, 2_{\bar{q}}; 5_\gamma, 6_\gamma \right) + B_{4;2\gamma}^0 \left(1_q, 4_g, 3_g, 2_{\bar{q}}; 5_\gamma, 6_\gamma \right) \right) \right] \right. \\
 &\quad \left. - \frac{\delta_{f_2}^{f_1}}{N_c^2} \frac{Q_f^4}{2} \tilde{B}_{4;2\gamma}^0 \left(1_q, 3_g, 4_g, 2_{\bar{q}}; 5_\gamma, 6_\gamma \right) \right]
 \end{aligned} \tag{3.5.36}$$

$$\begin{aligned}
& + \frac{\delta_{f_2}^{f_1}}{N_c} \left[n_u C_{4;2\gamma}^0 (1_{q_f}, 2_{\bar{q}_f}; 4_{\bar{u}}, 3_u; 5_\gamma, 6_\gamma) + n_d C_{4;2\gamma}^0 (1_{q_f}, 2_{\bar{q}_f}; 4_{\bar{d}}, 3_d; 5_\gamma, 6_\gamma) \right] \\
& + \frac{1}{N_c} C_{4;2\gamma}^0 (1_{q_{f_1}}, 3_{q_{f_1}}; 4_{\bar{q}_{f_2}}, 2_{\bar{q}_{f_2}}; 5_\gamma, 6_\gamma) \\
& - \frac{\delta_{f_2}^{f_1}}{N_c^2} D_{4;2\gamma}^0 (1_{q_f}, 3_{q_f}; 4_{\bar{q}_f}, 2_{\bar{q}_f}; 5_\gamma, 6_\gamma) \Big] J_0^2,
\end{aligned}$$

where n_u, n_d are the number of light up- and down-type quarks considered respectively,³ and represent the sum over the flavours of the final-state quark-antiquark pair created by the splitting of a gluon in the second Feynman diagram of fig. 3.2b. $C_{4;2\gamma}^0$ and $D_{4;2\gamma}^0$ both represent matrix-elements with four external quarks, and hence two pairs. Since a gluon connects the two quark pairs, for quark-antiquark scattering the Feynman amplitude is of the form

$$C_{4;2\gamma}^0 (q_1, \bar{q}_2, q_3, \bar{q}_4, \gamma, \gamma) \propto \delta_{f_2}^{f_1} \delta_{f_4}^{f_3} T_{i_2 i_1}^a T_{i_3 i_4}^a \mathcal{A}_s + \delta_{f_3}^{f_1} \delta_{f_4}^{f_2} T_{i_3 i_1}^a T_{i_2 i_4}^a \mathcal{A}_t \quad (3.5.37)$$

where $\mathcal{A}_s$ and $\mathcal{A}_t$ indicate amplitudes for s -channel quark-antiquark annihilation and t -channel gluon exchange respectively,⁴ the $\{f_j\}$ are the flavours of the quarks, and the $\{i_j\}$ their colours. The amplitudes $\mathcal{A}_s$ and $\mathcal{A}_t$ are related by a crossing symmetry,

$$\mathcal{A}_s = \mathcal{C}_{4;2\gamma}^0 (q_1, \bar{q}_2; \bar{q}_4, q_3; \gamma, \gamma) \quad (3.5.38)$$

$$\mathcal{A}_t = -\mathcal{C}_{4;2\gamma}^0 (q_1, q_3; \bar{q}_4, \bar{q}_2; \gamma, \gamma). \quad (3.5.39)$$

Squaring eq. (3.5.37), applying the Fierz identity eq. (2.3.38) and summing over colours gives

$$|C_{4;2\gamma}^0|^2 \propto (N_c^2 - 1) \left[\delta_{f_2}^{f_1} \delta_{f_4}^{f_3} |\mathcal{A}_s|^2 + \delta_{f_3}^{f_1} \delta_{f_4}^{f_2} |\mathcal{A}_t|^2 - \delta_{f_2}^{f_1} \delta_{f_3}^{f_1} \delta_{f_4}^{f_1} \frac{2}{N_c} \text{Re } \mathcal{A}_s^\dagger \mathcal{A}_t \right]. \quad (3.5.40)$$

Summed over final-state flavours, this gives

$$\sum_{f_3, f_4} |C_{4;2\gamma}^0|^2 \propto \left[\delta_{f_2}^{f_1} n_u |C_{4;2\gamma}^0 (q_1, \bar{q}_2; \bar{u}_4, u_3; \gamma, \gamma)|^2 \right] \quad (3.5.41)$$

³ Throughout we will take $n_u = 2$ and $n_d = 3$, excluding only the top; see table 1.1. In principle the top-quark contributions should be calculated separately in massive QCD table 1.1

⁴ For quark-quark scattering the same idea applies, but for gluon exchange in the t - and u -channels rather than annihilation in the s -channel.

$$\begin{aligned}
& + \delta_{f_2}^{f_1} n_d \left| \mathcal{C}_{4;2\gamma}^0 (q_1, \bar{q}_2; \bar{d}_4, d_3; \gamma, \gamma) \right|^2 \\
& + \left| \mathcal{C}_{4;2\gamma}^0 (q_1, q_3; \bar{q}_4, \bar{q}_2; \gamma, \gamma) \right|^2 \\
& - \frac{\delta_{f_2}^{f_1}}{N_c} \left(-2 \operatorname{Re} \left[\mathcal{C}_{4;2\gamma}^0 (q_1, \bar{q}_2; \bar{q}_4, q_3; \gamma, \gamma)^\dagger \mathcal{C}_{4;2\gamma}^0 (q_1, q_3; \bar{q}_4, \bar{q}_2; \gamma, \gamma) \right] \right).
\end{aligned}$$

which correspond term-by-term to the C - and D -type matrix elements of eq. (3.5.36).

Unlike the case of the two-quark B -type matrix elements, the electromagnetic charge associated with the photon emissions is not factorisable from within $\mathcal{C}_{4;2\gamma}^0$ and $D_{4;2\gamma}^0$ without further decomposition, and so they retain an implicit dependence upon the quark charges, as can be seen in the following decomposition of a constituent helicity amplitude (eq. (4.5) in [156]),

$$\begin{aligned}
\mathcal{C}_{4;2\gamma}^0 \left(1_{q_{f_1}}^-, 2_{\bar{q}_{f_1}}^+; 3_{q_{f_2}}^-, 4_{\bar{q}_{f_2}}^+; 5_\gamma^-, 6_\gamma^+ \right) &= Q_{f_1}^2 g_1(1, 2; 3, 4; 5, 6) + Q_{f_2}^2 g_1(3, 4; 1, 2; 5, 6) \\
&+ Q_{f_1} Q_{f_2} [g_2(1, 2; 3, 4; 5, 6) + g_2(3, 4; 1, 2; 5, 6)], \quad (3.5.42)
\end{aligned}$$

where the three terms correspond to double-emission from the first quark line, from the second, and one photon from each, respectively.

This complicates the construction of the subtraction term, since the electromagnetic factors must match in each limit if the divergences are to be correctly subtracted. This is a new problem, since simpler matrix elements have had the same charge factors as in the reduced matrix elements they map onto in divergent limits, so they could be extracted as a global factor.

To solve this we must keep track of quark-flavour. Formally, this can be achieved by introducing flavour-matched reduced matrix-elements $B_{2;2\gamma}^{0;f_i}$ and $B_{3;2\gamma}^{0;f_i}$ which are the corresponding matrix element weighted with the quark-charge of the i th quark, in the collinear limit in which the other quark pair maps onto a gluon,


$$B_{m;2\gamma}^{0;f_i} := Q_{f_i}^4 B_{m;2\gamma}^0. \quad (3.5.43)$$

With this we are able to write down the subtraction terms, for example:

$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{S NNLO}} = \left(8\pi^2\right)^2 N_{q\bar{q}}^{\text{LO}} \left(N_c^2 - 1\right) d\Phi_4 \quad (3.5.44)$$

$$\begin{aligned}
& \times \left[\delta_{f_2}^{f_1} \frac{Q_f^4}{2} \left[A_4^0(\hat{1}_q, 3_g, 4_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}) J_0^0 \right. \right. \\
& \quad + D_{3,q}^0(\hat{1}_q, 3_g, 4_g) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}'_q, 4'_g, \hat{2}'_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \\
& \quad + A_4^0(\hat{1}_q, 4_g, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}) J_0^0 \\
& \quad \left. + D_{3,q}^0(\hat{2}_q, 4_g, 3_g) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}'_q, 3'_g, \hat{2}'_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \right] \\
& - \frac{\delta_{f_2}^{f_1}}{N_c^2} \frac{Q_f^4}{2} \left[\tilde{A}_4^0(\hat{1}_q, 3_g, 4_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}) J_0^0 \right. \\
& \quad + A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}'_q, 4'_g, \hat{2}'_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \\
& \quad \left. + A_3^0(\hat{1}_q, 4_g, \hat{2}_{\bar{q}}) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}'_q, 3'_g, \hat{2}'_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \right] \\
& + \frac{\delta_{f_2}^{f_1}}{N_c} \left[n_u \left[B_4^0(\hat{1}_{q_f}, 3_u, 4_{\bar{u}}, \hat{2}_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1}(\hat{1}'_q, \hat{2}'_{\bar{q}}) J_0^0 \right. \right. \\
& \quad + \frac{1}{2} E_3^0(\hat{1}_{q_f}, 3_d, 4_{\bar{d}}) \left[B_{3;2\gamma}^{0:f_1} J_0^1 - A_3^0(\hat{1}'_{q_f}, 3'_g, \hat{2}'_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1} J_0^0 \right] \\
& \quad + \frac{1}{2} E_3^0(\hat{2}_{\bar{q}_f}, 3_d, 4_{\bar{d}}) \left[B_{3;2\gamma}^{0:f_1} J_0^1 - A_3^0(\hat{1}_{q_f}, 3'_g, \hat{2}'_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1} J_0^0 \right] \\
& \quad + n_d \left[B_4^0(\hat{1}_{q_f}, 3_d, 4_{\bar{d}}, \hat{2}_{\bar{q}_f}) B_{0;2\gamma}^{0:f_1}(\hat{1}'_q, \hat{2}'_{\bar{q}}) J_0^0 \right. \\
& \quad + \frac{1}{2} E_3^0(\hat{1}_{q_f}, 3_d, 4_{\bar{d}}) \left[B_{3;2\gamma}^{0:f_1} J_0^1 - A_3^0(\hat{1}'_{q_f}, 3'_g, \hat{2}'_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1} J_0^0 \right] \\
& \quad \left. + \frac{1}{2} E_3^0(\hat{2}_{\bar{q}_f}, 3_d, 4_{\bar{d}}) \left[B_{3;2\gamma}^{0:f_1} J_0^1 - A_3^0(\hat{1}_{q_f}, 3'_g, \hat{2}'_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1} J_0^0 \right] \right] \\
& + \frac{1}{N_c} \left[B_4^0(\hat{1}_{q_{f_1}}, 2_{\bar{Q}_{f_2}}, 3_{\bar{Q}_{f_2}}, 4_{q_{f_1}}) B_{0;2\gamma}^{0:f_1} J_0^0 \right. \\
& \quad + B_4^0(3_{\bar{Q}_{f_2}}, 4_{q_{f_1}}, 1_{q_{f_1}}, \hat{2}_{\bar{Q}_{f_2}}) B_{0;2\gamma}^{0:f_2} J_0^0 \\
& \quad - E_3^0(\hat{1}_{q_{f_1}}, \hat{2}_{\bar{Q}_{f_2}}, 3_{\bar{Q}_{f_2}}) \left[B_{1;2\gamma}^{0:f_1} J_0^1 + A_3^0(\hat{1}'_{q_{f_1}}, \hat{2}'_g, 4_{q_{f_1}}) B_{0;2\gamma}^{0:f_1} J_0^0 \right] \\
& \quad \left. - E_3^0(\hat{2}_{\bar{Q}_{f_2}}, \hat{1}_{q_{f_1}}, 4_{q_{f_1}}) \left[B_{1;2\gamma}^{0:f_2} J_0^1 + A_3^0(\hat{2}'_{\bar{Q}_{f_2}}, \hat{1}'_g, 3_{\bar{Q}_{f_2}}) B_{0;2\gamma}^{0:f_2} J_0^0 \right] \right] \\
& - 2 \frac{\delta_{f_2}^{f_1}}{N_c^2} \left[C_4^0(\hat{1}_{q_f}, 3_{\bar{q}_f}, 4_{q_f}, \hat{2}_{\bar{q}_f}) B_{2;2\gamma}^{0:f_1} J_0^0 + C_4^0(\hat{2}_{\bar{q}_f}, 4_{q_f}, 3_{\bar{q}_f}, \hat{1}_{q_f}) B_{2;2\gamma}^{0:f_1} J_0^0 \right],
\end{aligned}$$

where the terms have been organised to cancel directly the singularities of the corresponding terms in eq. (3.5.36), and momentum arguments have been omitted where they are clear (such as the momenta of the two photons).



(a) Illustrative Feynman diagram for one-loop matrix element $B_{3;2\gamma}^1$.

(b) Illustrative Feynman diagram for one-loop matrix element $\bar{B}_{3;2\gamma}^1$.

Figure 3.3: Example real-virtual diagrams contributing to $d\hat{\sigma}_{q\bar{q}}^{\text{RV}}$, depicted in the $q\bar{q}$ channel.

3.5.3.2. Real-virtual

The real-virtual matrix elements are subject to the same constraints as the real matrix elements of section 3.5.2; the diagrams must contain one quark-antiquark pair, a gluon, and the two final-state photons. The loop can either arise from a virtual gluon connecting to the quark line, or from a quark loop (in general of different flavour) created by a $g \rightarrow q\bar{q}$ splitting. These two possibilities are shown in figs. 3.3a and 3.3b respectively.

The possible diagrams that can contribute are further constrained by Furry's theorem,⁵ which guarantees that diagrams in which a single photon is emitted from a quark loop sum to zero, and by the fact that quark loops from which no photons are emitted are either tadpoles, which vanish in dimensional regularisation, or loop corrections to the gluon propagator, which were systematically accounted for in the renormalisation procedure of section 2.3.1.2 and contributed to the running coupling. Therefore only diagrams in which both photons are emitted from the quark loop contribute to the final result, as shown in fig. 3.3b. These can be trivially summed over the possible (light) flavours of the quark in the loop, to give the electromagnetic

⁵ Furry's theorem asserts that the vacuum expectation of an odd number of electromagnetic currents vanishes as a consequence of charge conjugation symmetry; each current $\bar{\psi}\gamma^\mu\psi$ brings a factor of (-1) under conjugation, and yet by symmetry the vacuum expectation value must be unchanged. This implies that it vanishes. Diagrammatically this can be understood as each diagram with a fermionic line in a given direction cancelling against its antifermionic counterpart.

charge factor

$$\overline{Q_f^2} = \sum_{f \in \{u, d, \dots\}} Q_f^2 = n_u \left(\frac{2}{3}\right)^2 + n_d \left(-\frac{1}{3}\right)^2 = \frac{11}{9}. \quad (3.5.45)$$

With this the real-virtual contribution to the cross-section in the $q\bar{q}$ -channel can be written as

$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{RV}} &= (8\pi^2) N_{q\bar{q}}^{\text{LO}} (N_c^2 - 1) \delta_{f_2}^{f_1} d\Phi_3 \\ &\times \left[Q_f^4 \left[B_{3;2\gamma}^1(1_q, 3_g, 2_{\bar{q}}; 4_\gamma, 5_\gamma) \right. \right. \\ &\quad - \frac{1}{N_c^2} \tilde{B}_{3;2\gamma}^1(1_q, 3_g, 2_{\bar{q}}; 4_\gamma, 5_\gamma) \\ &\quad \left. \left. + \frac{n_f}{N_c} \hat{B}_{3;2\gamma}^1(1_q, 3_g, 2_{\bar{q}}; 4_\gamma, 5_\gamma) \right] \right. \\ &\quad \left. + \overline{Q_f^2} \frac{1}{N_c} \overline{B}_{3;2\gamma}^1(1_q, 3_g, 2_{\bar{q}}; 4_\gamma, 5_\gamma) \right] J_0^1, \end{aligned} \quad (3.5.46)$$

where the contribution proportional to n_f , $\hat{B}_{3;2\gamma}^1$, arises purely from the renormalisation of the matrix element according to eq. (2.3.24) and so is simply proportional to $\varepsilon^{-1} B_{3;2\gamma}^0$,

$$\hat{B}_{3;2\gamma}^1 = -\frac{1}{\varepsilon} \left(-\frac{2}{3} T_F\right) B_{3;2\gamma}^0 = \frac{1}{3\varepsilon} B_{3;2\gamma}^0. \quad (3.5.47)$$

The corresponding subtraction terms $d\hat{\sigma}_{ab}^{\text{T}}$ arise both directly, using (unintegrated) antenna functions to cancel the singularities of the above matrix elements, and indirectly, including the (integrated) antenna functions from the double-real subtraction terms from eq. (3.5.44),

$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{T NNLO}} &= (8\pi^2) N_{q\bar{q}}^{\text{LO}} (N_c^2 - 1) d\Phi_3 \\ &\times \left[\delta_{f_2}^{f_1} Q_f^4 \left[A_3^1(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 + A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^1 \right] J_0^0 \right. \\ &\quad - \frac{1}{2} \mathcal{D}_{q,gg}^0(x_1) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \\ &\quad - \frac{1}{2} \mathcal{D}_{q,gg}^0(x_2) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \left. \right] \\ &- \frac{1}{N_c^2} \delta_{f_2}^{f_1} Q_f^4 \left[\tilde{A}_3^1(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 + A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^1 J_0^0 \right] \end{aligned} \quad (3.5.48)$$

$$\begin{aligned}
& - \mathcal{A}_{q\bar{q},g}^0(x_1, x_2) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \\
& + \frac{1}{N_c} \left[\mathcal{E}_{qQ,\bar{Q}}^0(x_1, x_2) \left[B_{3;2\gamma}^{0;f_1} J_0^1 + A_3^0(2_{\bar{q}}, 1_q, 3_g) B_{2;2\gamma}^{0;f_1} J_0^0 \right] \right. \\
& \quad + \mathcal{E}_{qQ,\bar{Q}}^0(x_2, x_1) \left[B_{3;2\gamma}^{0;f_2} J_0^1 + A_3^0(1_q, 2_{\bar{q}}, 3_g) B_{2;2\gamma}^{0;f_2} J_0^0 \right] \\
& \quad + \frac{n_f}{N_c} \delta_{f_2}^{f_1} Q_f^4 \left[\hat{A}_3^1(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right. \\
& \quad \quad - \frac{1}{2} \mathcal{E}_{qQ,\bar{Q}}^0(x_1, x_2) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \\
& \quad \quad \left. \left. - \frac{1}{2} \mathcal{E}_{qQ,\bar{Q}}^0(x_2, x_1) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0(\hat{1}_q, 3_g, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0 J_0^0 \right] \right] \right].
\end{aligned}$$


The integrated antenna functions here are the integrals of the $3 \rightarrow 2$, X_3^0 antenna functions in the double-real subtraction terms, accordingly mapped down from the double-real (6-parton) kinematics into the real-virtual (5-parton) kinematics, and can be compared at each color-level to those of eq. (3.5.44).

The remaining $4 \rightarrow 2$ double-real subtraction terms using X_4^0 antenna functions, and the $3 \rightarrow 2$ real-virtual subtraction terms using X_3^1 and X_3^0 antenna functions, map under integration into the double-virtual (4-parton) kinematics.


3.5.3.3. Double-virtual

The two-loop, four-parton, double-virtual matrix elements were calculated in [155], and the one-loop four-parton gluon-gluon box diagram was calculated in [159]. Illustrative diagrams contributing to the double-virtual contribution are shown in fig. 3.4.


Here we focus on the (integrated) subtraction terms in the $q\bar{q}$ -channel, $d\hat{\sigma}_{q\bar{q}}^U$. As explained above, since there are no additional particles that could become unresolved, there are no new unintegrated subtraction terms required; instead, we must include the image under integration of those double-real and real-virtual subtraction terms which map into the 4-parton kinematics, to ensure that the integrated and unintegrated subtraction terms cancel against each other. When combined with the appropriate mass-factorisation terms, the ε -poles of the integrated subtraction terms will cancel against the ε -poles of the two-loop matrix elements, as given by



(a) Two-loop diagram contributing to $B_{2;2\gamma}^2$, calculated in [155].



(b) Two-loop diagram contributing to $\bar{B}_{2;2\gamma}^2$, with gauge-theory factor proportional to $\bar{Q}_f^2$ and independent of $Q_{f1,2}^2$, calculated in [155].



(c) One-loop gluon-fusion 'box' diagram, first calculated in [159]. This diagram first contributes at NNLO for $p\bar{p} \rightarrow \gamma\gamma$ due to the relative suppression of the gluonic channels in the perturbative expansion; it is the leading-order contribution to $gg \rightarrow \gamma\gamma$.

Figure 3.4: Example four-parton loop Feynman diagrams contributing to the double-virtual corrections at NNLO.

the Catani singularity operators of eq. (2.3.33).

$$\begin{aligned}
d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{U NNLO}} = & N_{q\bar{q}}^{\text{LO}} (N_c^2 - 1) d\Phi_2 \\
& \times \left[\delta_{f_2}^{f_1} Q_f^4 \left[- \left[\mathcal{A}_{q\bar{q},gg}^0 + \mathcal{A}_{q\bar{q},g}^1 - b_0^{N_c} \mathcal{A}_{q\bar{q},g}^0 \log \frac{s_{12}}{\mu_R^2} \right] B_{2;2\gamma}^0 - \mathcal{A}_{q\bar{q},g}^0 B_{2;2\gamma}^1 \right] \right. \\
& - \frac{1}{N_c^2} \delta_{f_2}^{f_1} Q_f^4 \left[- \left[\frac{1}{2} \tilde{\mathcal{A}}_{q\bar{q},gg}^0 + \tilde{\mathcal{A}}_{q\bar{q},g}^1 + 2\mathcal{C}_{q\bar{q},q\bar{q}}^0 + 2\mathcal{C}_{q\bar{q},q\bar{q}}^0 \right] B_{2;2\gamma}^0 - \mathcal{A}_{q\bar{q},g}^0 B_{2;2\gamma}^1 \right] \\
& + \frac{1}{N_c} \left[- \mathcal{B}_{qQ,q\bar{Q}}^0 B_{2;2\gamma}^{0;f_1} - \mathcal{B}_{Qq,q\bar{Q}}^0 B_{2;2\gamma}^{0;f_2} \right] \\
& \left. + \frac{n_f}{N_c} \delta_{f_2}^{f_1} Q_f^4 \left[- \left[\mathcal{B}_{q\bar{q},Q\bar{Q}}^0 + \hat{\mathcal{A}}_{q\bar{q},g}^1 - b_0^{n_f} \mathcal{A}_{q\bar{q},g}^0 \log \frac{s_{12}}{\mu_R^2} \right] B_{2;2\gamma}^0 \right] \right],
\end{aligned} \tag{3.5.49}$$

where we have decomposed the first coefficient of the β -function of eq. (2.3.14) according to its group theory factors,

$$\beta_0 = \frac{11}{6} N_c - \frac{1}{3} n_f =: b_0^{N_c} N_c + b_0^{n_f} n_f. \tag{3.5.50}$$

The origins of each of these terms in the double-real or real-virtual subtraction counterterms can be inferred from the multiplicity of the antenna function and the loop-level, and can be systematically checked against them, noting that by definition,

$$n_f = n_u + n_d, \tag{3.5.51}$$

where the flavour distinction is made explicitly for the unintegrated antenna functions associated with a final-state quark pair, but not necessarily for the corresponding integrated subtraction terms.

3.5.3.4. Mass factorisation

Finally, we turn to the mass-factorisation counterterms. The double-virtual (4-parton) mass-factorisation counterterm includes the second-order coefficients of the mass-factorisation kernels for the first time. To clarify its structure, we follow [64, 77, 151] and define the *reduced* mass-factorisation kernels by

$$\bar{\Gamma}_{ab}^2(x) = -\frac{1}{2\varepsilon} \left[P_a^{b(1)}(x) + \frac{1}{\varepsilon} \beta_0 P_a^{b(0)}(x) \right], \tag{3.5.52}$$

so that the second-order convolution arising at NNLO simplifies,

$$\begin{aligned} \Gamma_{ca}^2(x_1) \delta_{db} \delta_{x_2}^1 + \Gamma_{ca}^1(x_1) \Gamma_{db}^1(x_2) + \delta_{x_1}^1 \delta_{ca} \Gamma_{db}^2(x_2) = \\ \bar{\Gamma}_{ca}^2(x_1) \delta_{db} \delta_{x_2}^1 + \delta_{x_1}^1 \delta_{db} \bar{\Gamma}_{ca}^2(x_2). \end{aligned} \quad (3.5.53)$$

The first-order, and second-order reduced, mass-factorisation kernels contain colour factors through their dependence upon the splitting functions of eqs. (2.1.7) to (2.1.10). To incorporate their colour-decomposition into the global composition (which is non-trivial, unlike in the NLO case), we define [64]

$$\Gamma_{qq}^1(x) = 2C_F \bar{\Gamma}_{qq}^1(x) \quad (3.5.54)$$

$$\Gamma_{gq}^1(x) = 2C_F \bar{\Gamma}_{gq}^1(x) \quad (3.5.55)$$

$$\Gamma_{qg}^1(x) = \bar{\Gamma}_{qg}^1(x) \quad (3.5.56)$$

$$\Gamma_{gg}^1(x) = N_c \bar{\Gamma}_{gg}^1(x) + n_f \hat{\Gamma}_{gg}^1(x) \quad (3.5.57)$$

for the first-order mass-factorisation kernels, and

$$\bar{\Gamma}_{qq}^2(x) = (N_c^2 - 1) \left[\bar{\Gamma}_{qq}^2(x) + \frac{1}{N_c} \tilde{\Gamma}_{qq}^2(x) + \frac{1}{N_c^2} \tilde{\tilde{\Gamma}}_{qq}^2(x) + \frac{n_f}{N_c} \hat{\Gamma}_{qq}^2(x) \right] \quad (3.5.58)$$

$$\bar{\Gamma}_{qQ}^2(x) = 2C_F \bar{\Gamma}_{qQ}^2(x) \quad (3.5.59)$$

for the two second-order mass-factorisation kernels that shall arise in the $q\bar{q}$ -channel.

The remainder are given, with consistent conventions, in eq. (A.20) of Ref. [64].

The real-virtual mass-factorisation counterterm is given by

$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{MF RV}} = -C_\varepsilon \sum_c \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \left[\delta_{x_1}^1 \left[d\hat{\sigma}_{q_{f_1}c}^{\text{R}} - d\hat{\sigma}_{q_{f_1}c}^{\text{S NLO}} \right] \Gamma_{c\bar{q}_{f_2}}^1(x_2) \right. \\ \left. + \Gamma_{cq_{f_1}}^1(x_1) \left[d\hat{\sigma}_{c\bar{q}_{f_2}}^{\text{R}} - d\hat{\sigma}_{c\bar{q}_{f_2}}^{\text{S NLO}} \right] \delta_{x_2}^1 \right] \end{aligned} \quad (3.5.60)$$

$$\begin{aligned} = - \left(8\pi^2 \right) N_{q\bar{q}}^{\text{LO}} (N_c^2 - 1) C_\varepsilon d\Phi_3 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \\ \times \left[\delta_{f_2}^{f_1} Q_f^4 \left[\left(\bar{\Gamma}_{qq}^1(x_1) \delta_{x_2}^1 + \delta_{x_1}^1 \bar{\Gamma}_{qq}^1(x_2) \right) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0 \left(\hat{1}_q, 3_g, \hat{2}_{\bar{q}} \right) B_{2;2\gamma}^0 J_0^0 \right] \right] \right. \\ \left. - \frac{1}{N_c^2} \delta_{f_2}^{f_1} Q_f^4 \left[\left(\bar{\Gamma}_{qq}^1(x_1) \delta_{x_2}^1 + \delta_{x_1}^1 \bar{\Gamma}_{qq}^1(x_2) \right) \left[B_{3;2\gamma}^0 J_0^1 - A_3^0 \left(\hat{1}_q, 3_g, \hat{2}_{\bar{q}} \right) B_{2;2\gamma}^0 J_0^0 \right] \right] \right] \end{aligned} \quad (3.5.61)$$

$$\begin{aligned}
& + \frac{1}{N_c} \frac{1}{1-\varepsilon} \left[\delta_{x_1}^1 \left[B_{3;2\gamma}^{0;f_1} J_0^1 + A_3^0 \left(\hat{1}_q, \hat{2}_g, 3_{\bar{q}} \right) B_{2;2\gamma}^{0;f_1} J_0^0 \right] \bar{\Gamma}_{gq}^1(x_2) \right. \\
& \quad \left. + \bar{\Gamma}_{gq}^1(x_1) \left[B_{3;2\gamma}^{0;f_2} J_0^1 + A_3^0 \left(\hat{2}_{\bar{q}}, \hat{1}_g, 3_q \right) B_{2;2\gamma}^{0;f_2} J_0^0 \right] \delta_{x_2}^1 \right],
\end{aligned}$$

where for brevity we have written

$$\delta_{x_i}^1 := \delta(1 - x_i), \quad (3.5.62)$$

and the incoming momentum arguments of the partonic cross-sections are

$$(\xi_1 x_1 P_1, \xi_2 x_2 P_2) \quad (3.5.63)$$

for partonic momentum fractions ξ_1 and ξ_2 respectively.

The $(1 - \varepsilon)^{-1}$ factor accompanying the final term accounts for the difference in spin-averaging factor between N_{qg} and $N_{q\bar{q}}$ within conventional dimensional regularisation (CDR); the number of helicities for quarks and gluons is

$$h_q^{\text{CDR}} = 2 \quad h_g^{\text{CDR}} = d - 2 = 2 - 2\varepsilon, \quad (3.5.64)$$

and so the averaging conversion factors are given by

$$N_{q\bar{q}} = \frac{h_g}{h_q} \cdot N_{qg} \quad N_{qg} = \frac{h_q}{h_g} \cdot N_{q\bar{q}} \quad (3.5.65)$$

with

$$\frac{h_g^{\text{CDR}}}{h_q^{\text{CDR}}} = 1 - \varepsilon \quad \frac{h_q^{\text{CDR}}}{h_g^{\text{CDR}}} = (1 - \varepsilon)^{-1}. \quad (3.5.66)$$

Identity-changing terms must also be corrected for the change in the number of possible initial-state colours, again through the averaging.

With the above conventions, the double-virtual (4-parton) mass-factorisation counterterm is given by

$$\begin{aligned}
d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{MF VV}} = & -C_\varepsilon^2 \sum_{c,d} \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \left[\Gamma_{cq_{f_1}}^1(x_1) \left[d\hat{\sigma}_{cd}^{\text{LO}} \right] \Gamma_{d\bar{q}_{f_2}}^1(x_2) \right. \\
& + \left[\delta(1 - x_1) \left[d\hat{\sigma}_{q_{f_1}c}^{\text{LO}} \right] \Gamma_{c\bar{q}_{f_2}}^2(x_2) \right. \\
& \quad \left. \left. + \Gamma_{c q_{f_1}}^2(x_1) \left[d\hat{\sigma}_{c\bar{q}_{f_2}}^{\text{LO}} \right] \delta(1 - x_2) \right] \right]
\end{aligned} \quad (3.5.67)$$

$$\begin{aligned}
& -C_\varepsilon \sum_c \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \left[\delta(1-x_1) \left[d\hat{\sigma}_{q_{f_1}c}^V - d\hat{\sigma}_{q_{f_1}c}^{T\text{ NLO}} \right] \Gamma_{c\bar{q}_{f_2}}^1(x_2) \right. \\
& \quad \left. + \Gamma_{cq_{f_1}}^1(x_1) \left[d\hat{\sigma}_{c\bar{q}_{f_2}}^V - d\hat{\sigma}_{c\bar{q}_{f_2}}^{T\text{ NLO}} \right] \delta(1-x_2) \right] \\
& = -N_{q\bar{q}}^{\text{LO}} (N_c^2 - 1) C_\varepsilon^2 d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \tag{3.5.68} \\
& \times \left[\delta_{f_2}^{f_1} Q_f^4 \left[\bar{\Gamma}_{qq}^1(x_1) \bar{\Gamma}_{qq}^1(x_2) + \delta_{x_1}^1 \bar{\Gamma}_{qq}^2(x_2) + \bar{\Gamma}_{qq}^2(x_1) \delta_{x_2}^1 \right. \right. \\
& \quad \left. \left. - \varepsilon^{-1} b_0^{N_c} \left[\delta_{x_1}^1 \bar{\Gamma}_{qq}^1 + \bar{\Gamma}_{qq}^1 \delta_{x_2}^1 \right] \right] B_{2;2\gamma}^0 \right. \\
& \quad \left. + \left[\delta_{x_1}^1 \bar{\Gamma}_{qq}^1(x_2) + \bar{\Gamma}_{qq}^1(x_1) \delta_{x_2}^1 \right] \left(B_{2;2\gamma}^1 + \mathcal{A}_{q\bar{q},g}^0 B_{2;2\gamma}^0 \right) \right] \\
& - \frac{1}{N_c^2} \delta_{f_2}^{f_1} Q_f^4 \left[\left[\bar{\Gamma}_{qq}^1(x_1) \bar{\Gamma}_{qq}^1(x_2) - \delta_{x_1}^1 \hat{\bar{\Gamma}}_{qq}^2(x_2) - \hat{\bar{\Gamma}}_{qq}^2(x_1) \delta_{x_2}^1 \right] B_{2;2\gamma}^0 \right. \\
& \quad \left. + \left[\delta_{x_1}^1 \bar{\Gamma}_{qq}^1(x_2) + \bar{\Gamma}_{qq}^1(x_1) \delta_{x_2}^1 \right] \left(B_{2;2\gamma}^1 + \mathcal{A}_{q\bar{q},g}^0 B_{2;2\gamma}^0 \right) \right] \\
& + \frac{1}{N_c} \left[\delta_{x_1}^1 \left[\bar{\Gamma}_{qQ}^2(x_2) - (1-\varepsilon)^{-1} \bar{\Gamma}_{gq}^1(x_2) \mathcal{A}_{qg,\bar{q}}^0 \right] B_{2;2\gamma}^{0;f_1} \right. \\
& \quad \left. + \delta_{x_2}^1 \left[\bar{\Gamma}_{qQ}^2(x_1) - (1-\varepsilon)^{-1} \bar{\Gamma}_{gq}^1(x_1) \mathcal{A}_{qg,\bar{q}}^0 \right] B_{2;2\gamma}^{0;f_2} \right] \\
& + \frac{n_f}{N_c} \delta_{f_2}^{f_1} Q_f^4 \left[\left[\delta_{x_1}^1 \hat{\bar{\Gamma}}_{qq}^2(x_2) + \hat{\bar{\Gamma}}_{qq}^2(x_1) \delta_{x_2}^1 - \varepsilon^{-1} b_0^{n_f} \left[\delta_{x_1}^1 \bar{\Gamma}_{qq}^1 + \bar{\Gamma}_{qq}^1 \delta_{x_2}^1 \right] \right] B_{2;2\gamma}^0 \right].
\end{aligned}$$

The mass-factorisation terms can be combined with the corresponding subtraction terms according to the colour level and the matrix element they contain, and systematically rearranged into ‘dipole’-like structures whose ε -poles explicitly mimic those of the Catani operators of eq. (2.3.33), to demonstrate that the poles do indeed cancel. The full details of this procedure are given in Ref. [64].

3.6. Implementation of diphoton production in NNLOJET

With the ingredients of the calculation identified, they can be implemented in a computer program to allow the numerical computation of cross-sections at the desired order of accuracy for a specific choice of physics parameters and fiducial cuts. This has been done in the NNLOJET framework [160].

NNLOJET is a parton-level Monte Carlo event generator that uses antenna

subtraction to compute cross-sections and differential cross-sections at fixed order, up to NNLO in QCD. It comprises around 3m lines of **Fortran** code and has to date been used for NNLO calculations of hadronic Z -boson production [161], $Z + j$ production [162], $W^\pm$ -boson production [163], $W^\pm + j$ production [164], $\gamma + j$ production [165], $\gamma\gamma$ production (this work), H production through vector-boson fusion [166], $H + j$ production through gluon fusion [167], $W^+ H + j$ production [168], single-jet production [169], and dijet production at leading colour [170], in addition to a number of lepton-lepton and hadron-lepton processes [171, 172].

The advantage to using a single framework for many processes is that ingredients universal to all fixed-order calculations can be implemented once, and reused many times. Within NNLOJET the code for the interface to the PDF library LHAPDF [173], the phase-space integration, and the analysis of the final-state (including fiducial cuts, jet clustering, photon isolation, which will be discussed shortly, and the differential binning of distributions into histograms) is shared between all processes in the ‘driver’ of the program, together with the necessary I/O infrastructure to allow the parameters of a calculation to be specified and the results extracted. For antenna subtraction, the fact that the subtraction is performed with a universal set of antenna functions and momentum maps also allows the code for these to be shared for all processes.

Each process has its own routines within the driver, programmatically generated from **Maple** scripts, which contain calls to each relevant matrix element, automatically dressed with the correct colour and symmetry factors, for each crossing in which they contribute.

A second set of **Maple** scripts takes **Maple** files encoding the subtraction terms for each such crossing of each matrix element, and automatically generates **Fortran** routines to compute them numerically, calling the shared code implementing the momentum maps and antenna functions where necessary.

This degree of automation and code-sharing removes much of the potential for manual error in the implementation of new processes in NNLOJET, but does not

remove it entirely.⁶ Since the purpose of investing effort into the complicated procedure of making higher-order predictions is to generate precise results suitable for comparison to data, such errors would risk undermining the whole enterprise. There is therefore an extensive series of independent checks performed on the component parts of each process to verify that they are correct, and on the result to check that their interplay has the expected properties. These are documented in appendix A.

⁶ This is set to be reduced still further with the emergence of performant automated one-loop provider codes such as OPENLOOPS [174] and BLACKHAT [175], which use numerical unitarity and recursion relations to construct tree-level and one-loop amplitudes, including those not yet amenable to analytical techniques.

Diphoton production: theory

Within perturbative QCD, the production of pairs of photons provides a final-state insensitive to the details of jet formation and hadronisation, and is therefore a good testing ground for our understanding of high-energy QCD. Beyond QCD, it provides a clean background against which to measure the properties of the Higgs boson [176,177], and as a possible channel for the detection of new physics [178,179].

Thus far we have focused exclusively on the perturbative QCD corrections to diphoton production, detailing the theoretical ingredients required, the obstacles, and a universal framework constructed to overcome them. The central question was how we might systematically regulate the divergences arising from the impossibility of separately resolving, for example, a gluon and a quark when they become arbitrarily collinear. Since $s_{qg} \rightarrow 0$ in this limit, any Feynman diagram with a quark propagator splitting to a collinear quark-gluon pair diverges $\propto s_{qg}^{-1}$ and must be combined with a virtual diagram to render the calculation finite.

Divergences of identical origin arise in $s_{q\gamma}$, from a quark propagator splitting to a final-state quark-photon pair. As in section 2.3.2.2, these are not physical divergences but artefacts of perturbation theory, and require additional theoretical machinery to properly account for them.

This requires the modelling of jet fragmentation into photons, in addition to

the ‘direct’ production of photons through the pointlike QED coupling in the hard-scattering. This is achieved through a final-state collinear factorisation equation, giving a convolution with ‘fragmentation functions’ analogous to the PDFs of chapter 2. This is described in section 4.1. The collinear quark-photon divergences arising from the pointlike QED coupling are then cancelled by a mass-factorisation counterterm in the fragmentation functions, exactly as for the collinear divergences associated with the PDFs in section 2.3.4.

In a hadron collider, photons are produced in abundance by fragmentation and the decay of final-state hadrons. To study photon processes, it is necessary to find ways to suppress this background of ‘indirect’ photons, relative to the desired signal of direct photons produced in the high-energy scattering. Since collinear fragmentation and hadronic decay both typically produce photons surrounded by a collimated spray of hadronic energy, one method is to identify photons as ‘isolated’ if they are not accompanied by hadronic energy. This is called an ‘isolation criterion’, introduced in section 4.2.

As a result, the fragmentation contribution is suppressed. If precisely-collinear hadronic radiation is vetoed by a given isolation criterion, the fragmentation contribution is suppressed entirely, and can simply be omitted from the calculation. This is the basis of ‘smooth-cone’ isolation, favoured by theorists. In practice, calorimeters cannot detect precisely-collinear radiation (only radiation falling into the same calorimeter cell), so such an isolation technique is impossible to implement experimentally. This mismatch between theory and experiment leads to an uncertainty in the predictions that is difficult to account for, as the theoretical isolation prescription must then itself simultaneously damp the collinear quark-photon singularity, and approximate the experimental isolation, to allow comparison of predictions with experiment.

In sections 4.2.2 and 4.2.3 we explore the theoretical properties of two alternatives: hybrid isolation, which attempts to compromise between the theoretical and the experimental isolations, and soft-drop isolation, a new method based on the inversion

of jet substructure techniques to identify photonic singletons within clustered jets.

In section 4.2.4 we will explore the effect of the discontinuity at the boundary of the isolated region of the photon that arises from the isolation veto on the predictions at NLO and NNLO, and in section 4.3 we will study the effect of the cuts on the perturbative convergence of the cross-section.

In preparation for applying the theory to phenomenology in chapter 5, where not otherwise specified we will use the fiducial cuts from the ATLAS 8 TeV experimental study [180],

$$\begin{aligned}
 p_{\text{T}}^{\gamma_1} &> 40 \text{ GeV} & p_{\text{T}}^{\gamma_2} &> 30 \text{ GeV} \\
 \Delta R_{\gamma\gamma} &> 0.4 & |y^\gamma| &\in [0, 1.37) \cup (1.56, 2.37) \\
 E_{\text{T}}^{\text{iso,part}} &< 11 \text{ GeV} & \text{within cone } \Delta R &\leq 0.4.
 \end{aligned} \tag{4.0.1}$$

4.1. Fragmentation

The collinear factorisation of eq. (2.1.3) arose from an attempt to distinguish between the long-distance, low-energy physics of a colliding hadron, and the short-distance high-energy collision. This led to parton distribution functions, satisfying perturbation-theory-derived evolution equations but requiring non-perturbative boundary conditions, and an overall convolution with the perturbatively-calculable partonic hard-scattering.

Since detected final-state photons are detected, they are intrinsically long-distance phenomena. This leads to a formal distinction between short-distance partonic photons, and resolved long-distance photons. The importance of such a distinction is clear for QCD, because confinement ensures that long-distance QCD partons can only be identified as jets, but must be imposed for photons.

An IR-safe way to define photon observables is therefore to strictly enforce this distinction through a factorisation relation, using an ‘ r -to- γ fragmentation function’¹

¹ Hadronisation, which also relates the partons produced in the final-state of the perturbative hard scattering to observable final-state particles, can be described in a similar way. This was the context

D_γ^r to account for the transition between a short-distance parton r produced by the hard scattering, and a long-distance observable photon γ . Photons are only considered observable once they have formally undergone this fragmentation process.

This gives rise to the overall factorisation relation

$$d\sigma_{AB \rightarrow X+\gamma} = \sum_{a,b} \sum_r f_a^A \otimes f_b^B \otimes d\hat{\sigma}_{ab \rightarrow X+r} * D_\gamma^r \quad (4.1.1)$$

where the convolution with the PDFs is as in eq. (2.1.3), and the final-state convolution $*$ is defined, with collinear four-momenta $p_\gamma = zp_r$, as

$$\left[d\hat{\sigma}_{ab \rightarrow X+r} * D_\gamma^r \right] (p_\gamma) = \int_0^1 \frac{dz}{z} d\hat{\sigma}(p_r) D_\gamma^r(z). \quad (4.1.2)$$

The sum over a and b here runs over the possible flavours of QCD partons as in eq. (2.1.3), whilst r may additionally be a photon itself.

Just as for initial-state collinear singularities and the PDFs in section 2.3.4, the ‘bare’ fragmentation function D_γ^{r0} must be adjusted with a mass-factorisation counterterm to compensate for the collinear singularities in the perturbative cross-section, to render the convolution finite. At leading-order, only the quark-photon splitting introduces a collinear singularity, so in the $\overline{\text{MS}}$ scheme [184] the corresponding renormalisation gives

$$D_\gamma^q(z, \mu_f^2) = D_\gamma^{0q}(z) - \frac{1}{\varepsilon} \left(\frac{4\pi\mu^2}{\mu_f^2} \right)^\varepsilon \frac{1}{\Gamma(1-\varepsilon)} P_\gamma^q(z), \quad (4.1.3)$$

where

$$P_\gamma^q(z) = Q_q^2 \frac{\alpha_{\text{em}}}{2\pi} \left[\frac{1 + (1-z)^2}{z} + \mathcal{O}(\alpha_{\text{em}}) \right]. \quad (4.1.4)$$

This leads to perturbative evolution equations for the dependence of the fragmentation functions on the fragmentation scale μ_f , entirely analogous to the DGLAP equations of eq. (2.1.6), save that the new possibility of $r = \gamma$ in eq. (4.1.1) introduces

in which such functions were first introduced by Feynman in [181], and developed into early QCD phenomenology [182, 183].

QCD-QED mixing,

$$\frac{\partial D_\gamma^r(z, \mu_f^2)}{\partial \ln \mu_f^2} = \sum_{s \in \{\gamma, q, \bar{q}, g\}} \sum_{m, n} \left(\frac{\alpha_s^m(\mu_f^2) \alpha_{\text{em}}^n(\mu_f^2)}{(2\pi)^{m+n}} \right) \int_z^1 \frac{dx}{x} P_s^{r(m, n)} \left(\frac{z}{x} \right) D_\gamma^s(x, \mu_f^2) \quad (4.1.5)$$

where to make contact with the single-expansion notation of section 2.1,

$$P_b^{a(m, 0)}(x) = P_b^{a(m-1)}(x). \quad (4.1.6)$$

The splitting functions arising in the evolution equations for final-state collinear factorisation are the ‘time-like’ splitting functions, identical to the space-like splitting functions of eqs. (2.1.7) to (2.1.10) at leading order and known up to NNLO [46, 185–187].

Since the photon is an elementary particle a partonic photon can become a long-distance photon by simply propagating, so

$$D_\gamma^r(x, \mu_f^2) = \delta(1-x) + \mathcal{O}(\alpha_{\text{em}}^2) \quad (4.1.7)$$

and so, neglecting the running of α_{em} ,

$$\begin{aligned} \frac{\partial D_\gamma^r(z, \mu_f^2)}{\partial \ln \mu_f^2} = \frac{\alpha_{\text{em}}}{2\pi} & \left[\sum_m \left(\frac{\alpha_s^m(\mu_f^2)}{(2\pi)^m} \right) P_\gamma^{r(m, 1)}(z) \right. \\ & \left. + \sum_{s \in \{q, \bar{q}, g\}} \sum_m \left(\frac{\alpha_s^m(\mu_f^2)}{(2\pi)^m} \right) \int_z^1 \frac{dx}{x} P_s^{r(m-1)} \left(\frac{z}{x} \right) D_\gamma^s(x, \mu_f^2) \right]. \end{aligned} \quad (4.1.8)$$

This inhomogeneous differential equation can only be fully solved with non-perturbative input to determine the boundary conditions for the solution of the complementary function. This input must be extracted from fits to data. For example, to leading order in the double expansion,

$$\frac{\partial D_\gamma^r(z, \mu_f^2)}{\partial \ln \mu_f^2} = Q_q^2 \frac{\alpha_{\text{em}}}{2\pi} P_\gamma^q(z) \quad (4.1.9)$$

and so

$$D_\gamma^q(z, \mu_f^2) = Q_q^2 \frac{\alpha_{\text{em}}}{2\pi} P_\gamma^q(z) \log \left(\frac{\mu_f^2}{\mu_0^2} \right) + D_\gamma^q(z, \mu_0^2). \quad (4.1.10)$$

The initial condition $D_\gamma^q(z, \mu_0^2)$ and the initial scale μ_0 must be determined from a

fit to data. This was first performed by the ALEPH collaboration at this order [188], and subsequently at NLO with data from electron-positron collisions at LEP [189,190] and later from DIS at HERA [191].

Exactly as for PDFs, the extraction of these functions from data introduces a considerable level of uncertainty, and dependence on the unphysical fragmentation scale. It is therefore desirable to find a way to make predictions that permit comparison with data but which are minimally exposed to the theoretical uncertainties of the fragmentation functions.

Inserting eq. (4.1.7) into eq. (4.1.1) gives a natural decomposition of the long-distance cross-section into a ‘direct’ contribution, in which an outgoing final-state photon becomes the resolved photon, and a ‘fragmentation’ contribution in which the fragmentation function is non-trivial and another parton r fragments,

$$d\sigma_{AB \rightarrow X+\gamma} = \sum_{a,b} f_a^A \otimes f_b^B \otimes \left[\underbrace{d\hat{\sigma}_{ab \rightarrow X+\gamma}}_{\text{‘direct’}} + \underbrace{\sum_{r \in \{q, \bar{q}, g\}} d\hat{\sigma}_{ab \rightarrow X+r} * D_\gamma^r}_{\text{‘fragmentation’}} \right]. \quad (4.1.11)$$

The ‘direct’ contribution is precisely what would be calculated in QCD if the final-state partonic photon and long-distance final-state photon were identified, whilst the fragmentation term arises from the distinction between the two.

To allow the calculation of the higher-order direct contribution to be tested experimentally without exposure to the fragmentation functions, it is desirable to define a fiducial region of phase-space in which the fragmentation contribution is heavily suppressed or exactly zero. There is then no error in approximating the long-distance cross-section $d\sigma_{AB \rightarrow X+\gamma}$ by the direct contribution alone. This is one of the central objectives of photon isolation, which will be introduced in section 4.2.

For diphoton production, there are two photons, and so a double convolution with fragmentation functions. When expanded in the same way as eq. (4.1.11), this gives ‘direct’, ‘single-fragmentation’ and ‘double-fragmentation’ contributions, examples of which are illustrated in fig. 4.1. These were calculated to NLO and implemented in the program DIPHOX [192], which remains the highest order of

fragmentation-inclusive diphoton production yet calculated.


4.2. Isolation

Photons produced by fragmentation are, by definition, accompanied by collinear QCD radiation. A straightforward way to isolate the direct production contribution and suppress that from fragmentation would therefore be to veto all events in which any QCD radiation is detected within a certain distance of the photon.


However, such a prescription fails to be infrared-safe, since the additional radiation of a soft gluon into the photon cone from any other particle would lead to the event being rejected. The theoretical implementation, meanwhile, would lead to a veto on the phase-space of soft-gluonic radiation, yet as seen in section 2.6 the full real-radiation phase-space must be integrated over to fully compensate for the ε -poles of the virtual matrix elements, which are unaffected by isolation. It is therefore necessary to allow some hadronic radiation in the neighbourhood of the photon. The way in which this is achieved is called a photon isolation prescription.

In addition to the theoretical motivation for the imposition of an isolation prescription, there is also a phenomenological motivation. Events at hadron colliders do not occur individually, and the collisions of the underlying event produce a background abundance of hadrons, many of which then decay to photon-pairs (such as η or π^0 mesons). Each such decay produces a highly-collimated photon pair, which is typically identified by a calorimeter as a single photon accompanied by hadronic radiation. For photonic final-states, these photons are produced in sufficient abundance to overwhelm the direct photon signal to which they form the background. To test QCD we must therefore additionally seek to suppress this background.


Despite early concerns that photon isolation would disrupt the factorisation of the cross-section of eq. (2.1.3) [193, 194], its validity for isolated cross-sections was proved in [195] and [196], provided that the corresponding measurement function F_{iso}^n satisfies infrared and collinear safety of the non-photon partons, and respects



(a) Leading order direct diphoton production.



(b) Single fragmentation ('Bremsstrahlung').



(c) Double fragmentation.

Figure 4.1: Examples of diphoton production: the leading order contributions of direct production, single fragmentation and double fragmentation respectively. Large shaded vertices represent the parton distribution functions from initial-state collinear factorisation, whilst small hatched vertices represent the non-trivial contributions of fragmentation functions in final-state collinear factorisation.

collinear photon-parton splittings,

$$F_{\text{iso}}^{n+1}(P_1, P_2; p_\gamma; p_1, \dots, p_{n+1}) \xrightarrow{p_\gamma \parallel p_i} F_{\text{iso}}^n(P_1, P_2; p_\gamma + p_i; p_1, \dots, p_n). \quad (4.2.1)$$

The approach primarily adopted for hadron collider detectors, ‘fixed-cone isolation’, is to consider a photon isolated if the total hadronic transverse energy deposited within a fixed cone of radius $R = \sqrt{\Delta\phi^2 + \Delta\eta^2}$ around photon i in the (η, ϕ) -plane, $E_{\text{T}}^{\text{had}}(R)$, is smaller than some threshold,

$$E_{\text{T}}^{\text{had}}(R) \leq E_{\text{T}}^{\text{iso}}(\gamma_i), \quad (4.2.2)$$

where we allow the threshold to vary between photons and events, typically as an affine function of the transverse energy of the photon $E_{\text{T}}^{\gamma_i}$,

$$E_{\text{T}}^{\text{iso}}(\gamma_i) := E_{\text{T}}^{\text{thr.}} + \varepsilon_\gamma E_{\text{T}}^{\gamma_i}. \quad (4.2.3)$$

This threshold is set by experiment on a case-by-case basis, differing between studies of different processes.

Just as for jet algorithms in fig. 2.5, the experimental isolation criterion must be applied to calorimeter cells, to reconstructed particles, or to both, whilst theory predictions apply cuts to simulated partons. For comparison to parton-level predictions, the experimental cuts are therefore unfolded using detector simulations to an approximately-equivalent fiducial cut on simulated partons [197]. An example of the difference between these scenarios, based on the cuts used by ATLAS in their 8 TeV diphoton study which will be featured heavily in chapter 5, is shown in fig. 4.2.

Any isolation criterion that completely or partially suppresses the collinear fragmentation component will also suppress the mass-factorisation term introduced in eq. (4.1.3) to regulate the collinear quark-photon singularities of the direct contribution. As a result, the dependence of a theoretical calculation on the isolation parameters will have both physical and unphysical origins. Understanding these is crucial if we are to judge how well we expect our calculations to describe data.

We therefore turn to introduce several of the isolation criteria that have been




Figure 4.2: Graphical representation of the internal structure of the photon isolation cone as applied at ATLAS experimentally in [180], and the standard-cone isolation to which it is unfolded. To scale in (η, ϕ) -space.

applied to theoretical calculations, some of which have also been implemented experimentally or applied to experimental data. Similarly to jet algorithms, these can be divided into cone-based criteria and clustering-based criteria.

4.2.1. Smooth-cone isolation

Smooth-cone isolation, or ‘Frixione isolation’, generalises fixed-cone isolation, which imposes a constant threshold on the total hadronic transverse energy deposited within the isolation cone, to a threshold function imposed on the radial profile of the cumulative hadronic transverse energy,

$$E_T^{\text{had}}(r) \leq E_T^{\text{iso}}(\gamma_i) \chi(r; R) \quad \forall r \leq R. \quad (4.2.4)$$

The function $\chi(r; R)$ may be chosen freely subject to the requirement that it vetoes exactly-collinear radiation, however soft, so that

$$\lim_{r \rightarrow 0} \chi(r; R) = 0. \quad (4.2.5)$$

It is typically additionally required to be continuous, monotonic, and such that $\chi(R; R) = 1$ on the boundary of the cone. The original choice of χ introduced

in [198],

$$\chi(r; R) = \left(\frac{1 - \cos r}{1 - \cos R} \right)^n \equiv \left(\frac{\sin \frac{1}{2}r}{\sin \frac{1}{2}R} \right)^{2n}, \quad (4.2.6)$$

has been widely adopted in isolation studies since. The other profile function commonly used in the literature is

$$\chi(r; R) = \left(\frac{r}{R} \right)^{2n}, \quad (4.2.7)$$

and for $R \leq \frac{\pi}{2}$, these are approximately equal, since

$$\left(\frac{1 - \cos r}{1 - \cos R} \right)^n = \left(\frac{r}{R} \right)^{2n} \left(1 + \frac{n}{12} (R^2 - r^2) + \mathcal{O} \left(\left(\frac{R}{2} \right)^4 \right) \right). \quad (4.2.8)$$

The finite granularity of the angular resolution of calorimeters makes this condition impossible to implement exactly at detectors, though a discretised version has been applied at the level of reconstructed particles at OPAL [199] and investigated for the LHC [200]. Other isolation procedures that can be implemented both theoretically and experimentally have recently been proposed, such as ‘soft-drop isolation’ [201], based on jet substructure techniques and related both to ‘democratic isolation’ [184] and to smooth-cone isolation in specific limits. These however have not yet been commonly adopted. As a result, all experimental measurements of final-states containing isolated photons so far performed at the LHC use fixed-cone isolation, whilst the majority of next-to-next-to-leading-order (NNLO) QCD predictions [202–204] use smooth-cone isolation.

It is therefore important to consider the important differences between fixed-cone and smooth-cone isolation. Within the profile function formalism, fixed-cone isolation corresponds to the constant profile function $\chi(r) \equiv 1$, which does not satisfy eq. (4.2.5), and so is not a legitimate smooth-cone choice of χ .

The constraint eq. (4.2.5) ensures that smooth-cone isolation entirely suppresses the fragmentation contribution, and so the collinear divergences from quark-photon splittings are suppressed by the resulting cut on the collinear phase-space rather than the mass-factorisation term of the fragmentation function. This leaves the

finiteness of the cross-section dependent on the isolation parameters, since in the $R \rightarrow 0$ limit, the $E_T^{\text{iso}} \rightarrow \infty$ limit or the $n \rightarrow -\infty$ limit the isolation is removed (save for the points of exact collinearity) and the collinear limit is restored to the cross-section.

Clearly this generates a residual dependence of the results of a calculation on these parameters. Variation over these parameters can be attempted analogously to scale variation to assess the magnitude of the resulting theoretical uncertainty. However, these parameters only describe a subspace of the function space of possible choices for χ . We can understand the impact of the choice of χ by considering subsets of phase-space. For any two isolation schemes with matching E_T^{iso} and R and profile functions $\chi_1(r)$ and $\chi_2(r)$, if

$$\chi_1(0) = \chi_2(0) \quad \text{and} \quad \chi_1(r) \leq \chi_2(r), \quad \forall r \leq R, \quad (4.2.9)$$

it follows that the permitted phase-space for the former is a subset of that for the latter, and so on physical grounds we expect that

$$d\sigma_1 \leq d\sigma_2. \quad (4.2.10)$$

4.2.1.1. Parameter dependence

Thanks to the form of the smooth-cone profile function of eq. (4.2.6) we can derive the leading dependence of the isolated cross-section on the parameters, by approximating the matrix elements by the splitting functions in the relevant limits and integrating over the cone.

For a collinear quark-photon splitting, the leading divergence of eq. (4.1.4) is proportional to

$$s_{q\gamma}^{-1} = \frac{1}{E_T^q E_T^\gamma \Delta R_{q\gamma}^2}, \quad (4.2.11)$$

so using

$$\frac{1}{E_i} d^3 p_i = E_T^i dE_T^i d\phi_i dy_i, \quad (4.2.12)$$

the integral over the three-momentum of a collinear quark inside the cone becomes, for $r(\phi, y) = \sqrt{\phi^2 + y^2}$,


$$\begin{aligned}\sigma_{\text{cone}}^q &\sim \int_{-R}^R d\phi \int_{-\sqrt{R^2-\phi^2}}^{\sqrt{R^2-\phi^2}} dy \int_0^{E_{\text{T}}^{\text{iso}} \chi(r)} dE_{\text{T}} \frac{E_{\text{T}}}{E_{\text{T}} r^2} \\ &= 2\pi E_{\text{T}}^{\text{iso}} \int_0^R dr \frac{\chi(r)}{r} \\ &= 2\pi E_{\text{T}}^{\text{iso}} \left([\chi(r) \log r]_{r=0}^R - \int_0^R dr \chi'(r) \log r \right)\end{aligned}\tag{4.2.13}$$

and so the condition eq. (4.2.5) is seen to be a necessary condition to tame the logarithmic behaviour of the integral over the collinear quark-photon singularity. Since the inclusive integral (in the absence of isolation) must be independent of R , the remainder of the cross-section retains a logarithmic dependence on R whatever the profile function applied within the isolation cone, visible in the first plot of fig. 4.3a. This is a known feature of narrow-cones in both smooth- and fixed-cone isolation and the possibility of resumming these logarithms was explored in [195]. For the profile function of eq. (4.2.7), eq. (4.2.13) gives, for $n > 0$,


$$\sigma_{\text{cone}}^q \sim E_{\text{T}}^{\text{iso}} \frac{1}{R^{2n}} \int_0^R dr r^{2n-1} = \frac{1}{2n} E_{\text{T}}^{\text{iso}}.\tag{4.2.14}$$

This relationship can be seen empirically in the NLO results from NNLOJET shown in fig. 4.3a, with a clear linear relationship in the second plot and a reciprocal relationship in the third. This reciprocal relationship means that the cross-section calculation is especially exposed to the collinear divergence for small values of n (e.g. the equivalent of $n = 0.1$ in [205]), and so will converge more slowly in a Monte Carlo simulation.

For a soft gluon, the real-emission matrix-element diverges according to eq. (2.3.40) as $1/(E_{\text{T}}^g)^2$, and has the isolation cut applied to it, whilst the subtraction term eq. (3.5.12) will cancel the divergence but have no isolation applied, since in the reduced momentum set there is no gluon in the isolation cone. As a result, for some



(a) Parameter-dependence of smooth-cone isolation for quark-photon collinear splittings.



(b) Parameter-dependence of smooth-cone isolation with a gluon in the photon cone.

Figure 4.3: The parameter-dependence of smooth-cone isolation, according to the identity of the parton isolated against. $\Delta\sigma$ is the difference between the cross-section with the specified isolation parameters and the baseline parameters $R = 0.4$, $n = 1$, $E_T^{\text{thr.}} = 11$ GeV (indicated by the grey lines) that will be returned to in chapter 5. The remaining cuts are those of the ATLAS experimental results that will be discussed in chapter 5.

scale $E_T^{\max}$ defining the boundary of the soft-approximation,

$$\begin{aligned}\sigma_{\text{cone}}^g &\sim \int_{-R}^R d\phi \int_{-\sqrt{R^2-\phi^2}}^{\sqrt{R^2-\phi^2}} dy \int_0^{E_T^{\max}} dE_T E_T \frac{\Theta[E_T^{\text{iso}} \chi(r) - E_T] - 1}{(E_T)^2} \\ &= \pi \left[R^2 \log \left(\frac{E_T^{\text{iso}}}{E_T^{\max}} \right) + 2 \int_0^R r \log \chi(r) dr \right],\end{aligned}\quad (4.2.15)$$

which for the profile function of eq. (4.2.7) gives

$$\sigma_{\text{cone}}^g \sim \pi R^2 \left[\log \left(\frac{E_T^{\text{iso}}}{E_T^{\max}} \right) - n \right]. \quad (4.2.16)$$

The dependence on $E_T^{\max}$ is expected to cancel against the complementary integral over the unapproximated matrix-element.

This result matches the intuition that for a soft gluon, there is nothing special about the cone around the photon, and so the cross-section is simply proportional to the area of the cone. The logarithmic dependence on E_T^{iso} and the linear (and negative) dependence on n can be seen in the numerical cross-sections computed by NNLOJET in fig. 4.3b.

Finally, since the profile function is unphysical, we repeat the derivations of eqs. (4.2.13) and (4.2.15) for a general smooth χ defined by a power series satisfying eq. (4.2.5),

$$\chi(r; R) = \sum_{k=1}^{\infty} a_k \left(\frac{r}{R} \right)^k, \quad (4.2.17)$$

where we note that the condition for the cone boundary,

$$\chi(R; R) = 1 \quad (4.2.18)$$

would imply

$$\sum_k a_k = 1. \quad (4.2.19)$$

For the quark-photon splitting,

$$\sigma_{\text{cone}}^q \sim E_T^{\text{iso}} \left[a_1 + \sum_{k=2}^{\infty} \frac{a_k}{k} \right],$$

where the imposition of $a_0 = 0$ to satisfy eq. (4.2.5) prevents a divergent logarithm from arising in the $r \rightarrow 0$ collinear limit. For the gluon case, eq. (4.2.15) gives (subject to the constraint $\sum_{n=1}^{\infty} a_{l+n}(r/R)^n < a_l$ for the series expansion of the logarithm),

$$\sigma_{\text{cone}}^g \sim R^2 \left[\log \left(a_l \frac{E_T^{\text{iso}}}{E_T^{\text{max}}} \right) - \frac{1}{2}l + 2 \sum_{m=1}^{\infty} \frac{1}{m+2} \frac{a_{l+m}}{a_l} + \mathcal{O}(R) \right], \quad (4.2.20)$$

$$\leq R^2 \left[\log \left(a_l \frac{E_T^{\text{iso}}}{E_T^{\text{max}}} \right) - \frac{1}{2}l + \frac{2}{3} \left(\frac{1}{a_l} - 1 \right) + \mathcal{O}(R) \right], \quad (4.2.21)$$

where here l is the index of the first non-zero coefficient $a_l > 0$, and in the second line we have used eq. (4.2.19). These results simplify to give eqs. (4.2.13) and (4.2.15) for the conventional choice of profile function eq. (4.2.7),

$$l = 2n, \quad a_k = \delta_{lk}. \quad (4.2.22)$$

These results show that the linear scaling $\sim E_T^{\text{iso}}$ in the quark case, and $\sim R^2$ in the gluon case, are general features of smooth-cone isolation with a smooth profile function, with only the coefficients of the scaling, and not the leading behaviour, dependent on the details of the profile function.

For fixed parameters, however, the coefficients matter, since they parametrise the theoretical uncertainty implicit in a making specific choice of profile function. Setting $l = 2$ and imposing $a_k = 0$ for $k > 4$ to restrict to quartic profile functions, and $a_k > 0$ to guarantee monotonicity, the scaling of σ_{cone}^q with E_T^{iso} varies by at most a factor of 2, with the upper bound given by the conventional profile function eq. (4.2.7) with $n = 1$ and the lower bound with $n = 2$. More generally, variation over polynomials with a_k non-zero only for $k_1 \leq k \leq k_2$ generates a factor of k_2/k_1 , with the upper-bound given by eq. (4.2.7) with $n = \frac{1}{2}k_1$ and the lower-bound given by eq. (4.2.7) with $n = \frac{1}{2}k_2$. This justifies the use of the profile function of eq. (4.2.7) in isolation studies, as when varied over n it bounds the space of alternatives between the two degrees.

As can be seen in fig. 4.3, for this conventional choice of profile function, the

quark contribution dominates the variation over isolation parameters. It is clear from eq. (4.2.20) that this is an artefact of this choice of χ , and that it is easy to choose pathological profile functions that generate extreme scaling behaviour by allowing l to become large.² In the limit of $l \rightarrow \infty$ with $a_l = 1$, $\chi_l \rightarrow \Theta[r - R]$ and gives the profile function corresponding to total isolation against any hadronic radiation (as discussed in section 4.2). As expected, $\sigma_{\text{cone}}^q \rightarrow 0$ whilst $\sigma_{\text{cone}}^g \rightarrow -\infty$ due to the miscancellation between the soft radiation and its subtraction term.

Finally, it is important to note that because of the ansatz eq. (4.2.17), the results of this section do not apply to non-analytic profile functions, including those that are discontinuous or have a discontinuous derivative. Such a profile function will arise shortly, in the context of ‘hybrid’ isolation.

4.2.2. Hybrid isolation

Hybrid isolation was introduced in [206] and describes a family of profile functions which interpolate between smooth-cone isolation with a given profile function, and fixed-cone isolation. This isolation procedure was applied to an NNLO calculation of photon production in association with a jet using NNLOJET in [165].

It can be formulated as smooth-cone isolation with the profile function

$$\chi^{\text{hyb}}(r; R_d, R) = \begin{cases} E_1 \chi(r; R_d) & r \in [0, R_d] \\ E_2 & r \in (R_d, R]. \end{cases} \quad (4.2.23)$$

As in eq. (4.2.3), E_1 and E_2 are, in general, affine functions of the photon transverse momenta. For $E_1 \leq E_2$, this is equivalent to applying fixed-cone isolation on the cone $r \leq R$ in addition to smooth-cone isolation on an inner cone $r \leq R_d$; for $E_1 > E_2$, these two formulations differ on the inner annulus $r \in (R_{\text{eff}}, R_d]$ on which $\chi(r; R_d) > E_2/E_1$. The latter formulation is then equivalent to a variant of the former, eq. (4.2.23), with a smaller effective radius $R_{\text{eff}} < R_d$. In the limit $R_d \rightarrow R$,

² The pathological behaviour in the $a_l \rightarrow 0$ limit is spurious, and arises from the breakdown of the expansion.

hybrid isolation reduces to smooth-cone isolation with the profile function χ , whilst the pointwise limit as $R_d \rightarrow 0$ corresponds to the fixed-cone profile function, except at $r = 0$, where the former is 0 and the latter 1.

From eqs. (4.2.10) and (4.2.23) we can deduce that the hybrid isolation cross-section grows as R_d decreases; intuitively, it grows as additional radiation is permitted within the isolation cone. Because the fragmentation contribution is vetoed by the value of the profile function at $r = 0$, the R_d parameter acts as the sole regulator of the collinear quark-photon singularity, independently of the constraint upon the radius of the outer photon cone R imposed by experimental isolation. The cross-section will therefore diverge logarithmically as $R_d \rightarrow 0$, as the collinear singularity is exposed. It follows that there is some value of the parameter R_d for which the hybrid cross-section and the fixed-cone cross-section must coincide, and the divergent cross-section of vetoed radiation in the inner-cone numerically matches that of the missing fragmentation counterterm. In section 5.1 we will apply this principle to ensure that our choice of R_d for phenomenology is reasonable.

With a view to applying hybrid isolation to diphoton phenomenology, we will focus on its implications for cross-sections and differential distributions, both through its dependence on the inner-cone radius R_d , and in comparison to conventional smooth-cone isolation.

4.2.2.1. Parameter dependence

As in section 4.2.1.1 we can apply the factorisation relations of matrix elements in the parton-in-cone limit to derive the (single-parton-in-cone) dependence of the cross-section on the isolation parameters.

We will assume that $E_1 \leq E_2$. The inner-cone is then exactly as in eqs. (4.2.13) and (4.2.15), whilst the integral over the standard-cone outer-annulus can be obtained by changing the integration region to $r \in [R_d, R]$ and setting $\chi(r) \equiv 1$,

$$\sigma_{\text{cone}}^q \sim \frac{E_1}{2n} - E_2 \log \frac{R_d}{R} \quad (4.2.24)$$




Figure 4.4: The matched-hybrid isolation profile function χ^{hyb} for $n \in \{\frac{1}{2}, 1, 2\}$ and several choices of the inner cone radius, R_d (dashed). As $R_d \rightarrow 0$ (dotted), the smooth-cone (solid) suppression of the collinear singularity is retained, but the numerical deviation from the constant profile function of fixed-cone isolation is diminished. For all values of R_d , exactly-collinear radiation is vetoed.

$$\sigma_{\text{cone}}^g \sim R_d^2 \left[\log \frac{E_1}{E_2} - n \right] + R^2 \log \frac{E_2}{E_T^{\text{max}}}. \quad (4.2.25)$$

In the $R_d \rightarrow R$ limit the smooth-cone results of eqs. (4.2.13) and (4.2.15) are restored.

4.2.2.2. Matched-hybrid isolation

In chapter 5 we will chiefly consider the phenomenology of matched-hybrid isolation, where we impose continuity at the boundary between the inner-cone and the outer annulus: $E_1 = E_2$. Other choices are discontinuous at $r = R_d$, which is expected to lead to instabilities.³ In this scheme, when making experimental predictions, once the inner-cone profile function χ is chosen, the parameters E_T^{iso} and R are fixed by the fiducial cuts of the experiment. The only remaining unphysical parameter is then R_d , the radius of the inner cone.

³ For matched-hybrid isolation, only the derivative χ' is discontinuous at $r = R_d$. It is possible to define more sophisticated piecewise schemes which are arbitrarily smooth at R_d , and non-piecewise smooth-cone profile functions with similar properties to hybrid isolation, such as $\chi(r; R_d, R) = \left(1 + \exp \frac{r - R_d}{r(R - R_d)}\right)^{-1}$.

Parameter dependence of cross-sections

Since we are concerned with the physical criterion in eq. (4.2.10), we consider the hybrid-isolation cross-section relative to the corresponding smooth-cone prediction,

$$\Delta\sigma(R_d) = \sigma_{\text{hybrid}} - \sigma_{\text{smooth}}, \quad (4.2.26)$$

using the profile function of eq. (4.2.6). We can consider $\Delta\sigma(R_d)$ as the physical cross-section resulting from the presence of the generalised isolation measurement function

$$\Theta \left[\chi^{\text{hyb}}(\{p_i\}; R_d) - E_{\text{T}}^{\text{had}}(R) \right] - \Theta \left[\chi^{\text{smooth}}(\{p_i\}) - E_{\text{T}}^{\text{had}}(R) \right]. \quad (4.2.27)$$

in the integrand. This is zero for, and hence vetoes, events that are treated commonly by the two isolation criteria, and since

$$\chi^{\text{hyb}}(r; R_d, R) \geq \chi^{\text{smooth}}(r; R), \quad (4.2.28)$$

selects those that are vetoed under smooth-cone isolation but permitted under hybrid isolation. The Heaviside step functions implementing the isolation criteria induce discontinuities in the resulting distributions, which will be discussed further in section 4.2.4.

We begin by summarising the R_d -dependence of $\Delta\sigma(R_d)$, where other parameters are fixed, so R and $E_{\text{T}}^{\text{thr.}}$ are common to both profile functions. Where a gluon is emitted inside the cone, eq. (4.2.25) implies that

$$\Delta\sigma(R_d) \sim -nR_d^2, \quad (4.2.29)$$

in accordance with the intuition that the additional cross-section allowed is proportional to the area over which the gluon can additionally be emitted. Where a quark is emitted, eq. (4.2.24) gives

$$\Delta\sigma(R_d) \sim -\log \frac{R_d}{R}. \quad (4.2.30)$$




Figure 4.5: The variation $\Delta\sigma(R_d) = \sigma_{\text{hybrid}} - \sigma_{\text{smooth}}$ at NLO as a function of the inner-cone radius R_d , for $E_T^{\text{thr.}} = 11 \text{ GeV}$ and $R = 0.4$. The lines plotted in the second and third plots are simple logarithmic and quadratic fits, respectively; those in the first are their sum.

This behaviour is verified empirically at NLO, using NNLOJET, in fig. 4.5.

The dependence of the inner-smooth-cone cross-section on its remaining isolation parameters is that of eqs. (4.2.24) and (4.2.25).

Isolation effects and phase-space

At NLO the underlying kinematics restrict the relevance of photon isolation to a relatively minor region of phase-space. The only part of the fixed-order NLO calculation sensitive to the isolation parameters is the real emission, and within the real contribution, the final state parton p_1 may only enter the isolation cone of the second-hardest photon, as they must together balance $\mathbf{p}_T^{\gamma_1}$. The collinear invariant being regulated by the isolation criterion is therefore

$$s_{\gamma_2 p_1} \approx E_T^{\gamma_2} E_T^1 \Delta R_{\gamma_2 p_1}^2 = E_T^{\gamma_2} p_T^{\gamma\gamma} \Delta R_{\gamma_2 p_1}^2, \quad (4.2.31)$$

where

$$p_T^{\gamma\gamma} = \|\mathbf{p}_T^{\gamma_1} + \mathbf{p}_T^{\gamma_2}\| \quad (4.2.32)$$

is the transverse momentum of the diphoton system, and the last equality is valid only for three-particle final-states.

For any monotonic profile function χ , it follows from eq. (4.2.10) that the resulting isolation criterion is at least as restrictive as fixed-cone isolation with the same boundary condition, so the effect of isolation will be confined to $p_T^{\gamma\gamma} = E_T^1 \leq E_T^{\text{iso}}(\gamma_2)$ purely from kinematic constraints.⁴ This implies that any differences between two isolation schemes are only resolved at this order on the strip

$$p_T^{\gamma_2} \in \left[\max \left\{ p_T^{\gamma_2, \text{cut}}, \frac{p_T^{\gamma_1} - E_T^{\text{thr.}}}{1 + \varepsilon_\gamma} \right\}, p_T^{\gamma_1} \right]. \quad (4.2.33)$$

This is illustrated in practice in fig. 4.6, for $E_T^{\text{iso}} = 11$ GeV.

For asymmetric photon cuts with a p_T^{cut} -gap greater than $E_T^{\text{thr.}}$, this would exclude events close to the threshold of the photon cuts from isolation dependence entirely, at this order. For the more conventional case, the dependence of the NLO cross-section on the parameters is dominated by events on the threshold of the cuts.

Parameter dependence of differential cross-sections

The differential formulation of the isolation criteria in eqs. (4.2.13) and (4.2.15) is

$$\frac{d^3 \sigma_{\text{cone}}^{q/g}}{dE_T dy d\phi} \propto \Theta \left[E_T^{\text{iso}} \chi(r) - E_T \right], \quad (4.2.34)$$

which makes clear that the scaling of the cross-section of eq. (4.2.15) will not be distributed uniformly across all differential cross-sections. Instead, the variation in $\Delta\sigma$ shown above generated by variation of isolation parameters will be concentrated in certain regions of distributions, whilst others will be relatively insensitive to them. As discussed above, in the real-emission kinematics, the transverse momentum of the parton must balance against that of the diphoton system, $p_T^{\gamma\gamma} = E_T^1$, and so the Θ -function applies to $p_T^{\gamma\gamma}$ directly,

$$\Theta \left[E_T^{\text{iso}} \chi(r) - E_T \right] \equiv \Theta \left[E_T^{\text{iso}} \chi(r) - p_T^{\gamma\gamma} \right]. \quad (4.2.35)$$

⁴ As a consequence, for fixed radius R we would expect the constraints imposed by unitarity to force a larger choice of R_d for more restrictive isolation thresholds $E_T^{\text{thr.}}$, and to permit a smaller choice for less strict threshold energies.




Figure 4.6: The region of $(p_T^{\gamma_1}, p_T^{\gamma_2})$ -space on which differences between isolation prescriptions can be resolved for one and two real emissions (R- and RR-kinematics respectively). In the Born kinematics, isolation has no effect. The relevant cuts are $p_T^{\gamma_1} \geq 40$ GeV, $p_T^{\gamma_2} \geq 30$ GeV and $E_T^{\text{thr.}} = 11$ GeV, $\varepsilon_\gamma = 0$. The shaded bins are those populated in practice by 50m VEGAS phase-space points within NNLOJET, when adapted to the matrix elements, illustrating the dominance of the cross-section by events close to the threshold of the cuts.

As a result, the $d\Delta\sigma/dp_T^{\gamma\gamma}$ distribution will be zero above $p_T^{\gamma\gamma} = E_T^{\text{thr.}}$, for $\varepsilon_\gamma = 0$, and non-zero below it. We will investigate the consequences of this further in section 4.2.4. Because the entire variation of the cross-section with the isolation parameters occurs in this region, we can expect the isolation uncertainties in it and in the corresponding regions of correlated distributions to be very large, whilst regions of distributions that are predominantly populated by events of $p_T^{\gamma\gamma} > E_T^{\text{thr.}}$ will have much reduced sensitivity to the isolation parameters.

To explore the consequences of this further, in figs. 4.7 and 4.8 we show a selection of differential cross-sections $d\Delta\sigma(R_d)/dX$ for a range of values for R_d . These are the distributional counterparts to fig. 4.5. As in fig. 4.5, the cuts are chosen to match those used in the 8 TeV ATLAS study, whilst the theory parameter R_d is varied.




Figure 4.7: Isolation cone effects at NLO, showing the difference between matched-hybrid and smooth-cone isolation $\Delta\sigma$. The $d\Delta\sigma/d\Delta R_{\gamma j}$ distribution has regions of highly-local sensitivity to R_d , whilst the $d\Delta\sigma/d\Delta y_{\gamma\gamma}$ distribution is sensitive only through a small global normalisation. In the first plot the jet cut is 1 GeV; at this order all jets comprise a single parton. Higher values of the jet cut increase the minimal value of $\Delta R_{\gamma j}$ at which partons of that p_T can be emitted and not vetoed, with the minimum approximately $R\sqrt{p_T^j/E_T^{\text{iso}}}$, leading to steeper slopes to the left of the peak.

As can be seen in the first plot of fig. 4.7, the characteristic kinematic configuration of the events additionally allowed by hybrid isolation is very sensitive to the choice of R_d . The peak at R_d arises because the difference between the smooth-cone and hybrid profile functions is maximised at R_d . This leads to a localised sensitivity to the R_d parameter in certain distributions. This exposure of the collinear singularity shown in fig. 4.7 with decreasing R_d illustrates the kinematics underlying the logarithmic behaviour of eq. (4.2.30), and shows a gradual bias within the photon-cone towards increasingly collinear events as the inner-cone is reduced in size. In other distributions such as $d\Delta\sigma/d\Delta y_{\gamma\gamma}$, also shown in fig. 4.7, the logarithmic behaviour manifests itself only as a global normalisation.

Further distributions in which the effect is localised are shown in fig. 4.8 alongside the corresponding smooth-cone distributions. These illustrate interesting features of the isolated differential cross-sections at NLO. In the first figure, the $d\Delta\sigma/dp_T^{\gamma\gamma}$ distribution shows a discontinuity at $p_T^{\gamma\gamma} = E_T^{\text{iso}}$. The shape of this distribution is sensitive to the parameters of hybrid isolation and the offset between asymmetric photon cuts. Here, the peak occurs at the offset whilst the discontinuity occurs at E_2 , in the notation of eq. (4.2.23) (including for non-matched isolation). If E_2 were allowed to depend on $p_T^{\gamma_2}$ this discontinuity would be smoothed over an interval in $p_T^{\gamma\gamma}$, but would reappear in another distribution. This arises directly from the boundary of the fixed-cone criterion in phase-space and will be discussed further, including its consequences for higher-orders, in section 4.2.4.

The $d\Delta\sigma/dp_T^{\gamma_2}$ distribution, and as a direct consequence, the $d\Delta\sigma/dM_{\gamma\gamma}$ distribution, show discontinuities, in the differential cross-section and its derivative respectively, at the boundaries of the Born phase-space. The latter was analysed in [207]. The former arises because real soft QCD radiation is kinematically restricted to arise only close to the back-to-back configuration $p_T^{\gamma_2} \lesssim p_T^{\gamma_1}$, which is permitted by the isolation criteria by design, and cannot cancel as anticipated against virtual poles outside the Born kinematics.⁵

⁵ The kinematic prohibition of these soft emissions is the underlying mechanism for the unphysical

These (unphysical) features arise commonly in both smooth-cone and fixed-cone isolation. They are a direct consequence of the requirement that soft gluon radiation be permitted, to allow the general cancellation of real and virtual singularities. Where the virtual singularities are kinematically prohibited, but real soft singularities are not, a miscancellation arises.

Since the behaviour of the isolated cross-section at NLO is highly sensitive to the unphysical behaviour in these regions, it is *a priori* unclear to what extent the variation of isolation parameters based on NLO behaviour will lead to conclusions that hold at higher orders. Running enough calculations at NNLO with sufficient resolution to investigate the R_d -dependence of distributions in the regions of non-analyticity shown in fig. 4.8 would be prohibitively computationally expensive. In section 5.1.2 we will therefore compare smooth-cone isolation to matched-hybrid isolation with fixed $R_d = 0.1$.

To illustrate the overall dependence of the NNLO cross-section on R_d , in fig. 4.9 we show the NNLO counterpart to fig. 4.5. The dependence on R_d is again dominated by the qg channel. Overall, the magnitude of the effect is similar to that at NLO despite the contribution from events outside the strip of eq. (4.2.33), whilst the shape is no longer logarithmic. Channels in which a parton is permitted to enter the photon cone for the first time at NNLO have the same R_d -dependence shown in fig. 4.5. This suggests that the procedure used to justify the choice $R_d = 0.1$ above, by comparison to the fragmentation calculation, should remain valid at NNLO. However, it also suggests that the dependence of the cross-section on R_d at NNLO is dominated by the NLO real-radiation effects that ultimately give rise to unphysical behaviour. We will investigate the role of these features at NNLO in section 4.2.4.

dependence of σ_{NLO} on the p_T^γ cuts when moving from asymmetric to symmetric cuts, first remarked upon in the context of jet production in [208]. This can clearly be seen from the lower-right plot in fig. 4.8, and will be commented upon further in section 4.3.




Figure 4.8: Detailed isolation cone effects at NLO, showing the difference between matched-hybrid and smooth-cone isolation $\Delta\sigma$. The absolute predictions for smooth-cone isolation are shown for reference. At this order, isolation criteria only apply at all in the limited region of phase-space defined by $p_T^{\gamma\gamma} \leq E_T^{\text{iso}}$. Here, as for the ATLAS 8 TeV data considered throughout, $E_T^{\text{iso}} = 11$ GeV.




Figure 4.9: The variation $\Delta\sigma(R_d) = \sigma_{\text{hybrid}} - \sigma_{\text{smooth}}$ at NNLO as a function of the inner-cone radius R_d , for $R = 0.4$. All other parameters are kept constant. The two channels not shown, $q\bar{q}$ and gg , have comparable shape (but smaller magnitude) to qg and qq respectively.

4.2.3. Democratic isolation

4.2.3.1. Traditional democratic isolation

‘Democratic’ isolation was introduced in [184] as an alternative to cone-based isolation, with a conceptually different approach. In cone-based isolation, the photons are first identified, then the isolation criterion is applied to partonic radiation within their photon cones, and if the event passes the isolation criterion, the hadronic radiation is clustered into jets. In traditional democratic isolation, all final-state radiation, whether hadronic or electromagnetic, is clustered ‘democratically’ into jets; these jets are then classified as ‘photonic’ if the electromagnetic energy fraction,

$$z = \frac{E_{\text{em}}^j}{E^j}, \quad (4.2.36)$$

exceeds some threshold z_{cut} , and as hadronic if not.

Unlike smooth-cone isolation, democratic isolation retains its dependence upon the fragmentation contribution of eq. (4.1.11), which was exploited in [189, 191] to extract fits of the fragmentation functions from LEP and HERA data.

4.2.3.2. Soft-drop isolation

Soft-drop isolation was introduced in [201] as a new democratic isolation criterion inspired by jet substructure techniques.

As outlined in section 2.5.3, in conventional soft-drop [110], jets are clustered according to an initial generalised k_T -algorithm (typically anti- k_T), and the constituent partons of each jet are reclustered according to a second algorithm (conventionally Cambridge–Aachen). The resulting sequence of clusterings is then tested in reverse order, iteratively according to the harder of the two clustered subjects, until either the condition of eq. (2.5.6),

$$\frac{p_T^2}{p_T^1 + p_T^2} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R} \right)^\beta, \quad (4.2.37)$$

is met for the subjects $p_T^1 \geq p_T^2$, or only one particle remains. In the latter case, the jet can either be discarded (‘tagging’ mode), or returned as a singleton jet.

In soft-drop isolation, if the singleton jet comprises a photon, it is classified as an isolated photon.

This has a natural relationship with the traditional democratic isolation of section 4.2.3.1, both because the initial clustering treats hadronic and electromagnetic equally, and because the soft-drop condition with $\beta = 0$ is equivalent to the democratic isolation criterion for an identified photon eq. (4.2.36), applied to the reconstructed subjects within a jet. Its application to reclustered subjects rather than at the level of the jet ensures that soft-drop isolation requires focused, rather than diffuse, electromagnetic energy within the jet to identify it as photonic.

4.2.3.3. Parameter dependence

As in section 4.2.1.1 and section 4.2.2.1, we once again apply the splitting functions to the derivation of the parameter-dependence of soft-drop-isolated cross-sections, which we will verify against the NNLOJET implementation. Details of the implementation and validation of soft-drop isolation in NNLOJET can be found in appendix A.6.

For the photon to be identified as a photon, and not simply a constituent of a jet, it must be the harder of the two particles, and the soft-drop condition eq. (2.5.6) must have failed; if the condition had passed, the photon would not have been returned as a singleton and would instead have become part of a jet with the QCD parton. Therefore, assuming $\beta > 0$ and $z_{\text{cut}} < 1$,

$$p_{\text{T}}^q \leq \frac{p_{\text{T}}^\gamma}{z_{\text{cut}}^{-1} \left(\frac{\Delta R_{12}}{R} \right)^{-\beta} - 1} \quad (4.2.38)$$

$$=: p_{\text{T}}^\gamma \chi_{\text{SD}}(z_{\text{cut}}, \Delta R_{12}; R), \quad (4.2.39)$$

where we have defined the effective profile function χ_{SD} . This effective profile function can be seen to satisfy eq. (4.2.5), and so at NLO is exactly equivalent to smooth-cone isolation with this choice of profile function, despite being defined as a substructure isolation technique rather than a cone-based one. Indeed, for small z_{cut} ,

$$\chi_{\text{SD}}(z_{\text{cut}}, r; R) = \frac{z_{\text{cut}} p_{\text{T}}^\gamma \left(\frac{r}{R} \right)^\beta}{1 - z_{\text{cut}} \left(\frac{r}{R} \right)^\beta} \quad (4.2.40)$$

$$= z_{\text{cut}} p_{\text{T}}^\gamma \left(\frac{r}{R} \right)^\beta (1 + \mathcal{O}(z_{\text{cut}})), \quad (4.2.41)$$

and the effective soft-drop profile function coincides with that of smooth-cone isolation eq. (4.2.7), for

$$E_{\text{T}}^{\text{iso}} = z_{\text{cut}} p_{\text{T}}^\gamma \quad \beta = 2n, \quad (4.2.42)$$

in this limit.

We can therefore directly apply the result of eq. (4.2.13), to give

$$\sigma_{\text{cone}}^q \sim \int_{-R}^R d\phi \int_{-\sqrt{R^2-\phi^2}}^{\sqrt{R^2-\phi^2}} dy \int_0^{p_{\text{T}}^\gamma \chi_{\text{SD}}(r)} dE_{\text{T}} \frac{1}{r^2} \quad (4.2.43)$$

$$= 2\pi p_{\text{T}}^\gamma \int_0^R dr \frac{\chi_{\text{SD}}(r)}{r} \quad (4.2.44)$$

$$= -2\pi p_{\text{T}}^\gamma \frac{1}{\beta} \log[1 - z_{\text{cut}}] \quad (4.2.45)$$

$$\propto \frac{z_{\text{cut}} p_{\text{T}}^\gamma}{\beta} (1 + \mathcal{O}(z_{\text{cut}})). \quad (4.2.46)$$

This scaling can be seen to hold empirically in the plots of cross-sections obtained

by NNLOJET in fig. 4.10b.

Likewise, for the gluon-in-cone case, eq. (4.2.15) implies that

$$\begin{aligned}\sigma_{\text{cone}}^g &\sim \int_{-R}^R d\phi \int_{-\sqrt{R^2-\phi^2}}^{\sqrt{R^2-\phi^2}} dy \int_0^{E_T^{\text{max}}} dE_T E_T \frac{\Theta[p_T^\gamma \chi_{\text{SD}}(r) - E_T] - 1}{(E_T)^2} \\ &= \pi \left[R^2 \log \left(\frac{z_{\text{cut}} p_T^\gamma}{E_T^{\text{max}}} \right) + 2 \int_0^R r \log \chi_{\text{SD}}(r) dr \right],\end{aligned}\quad (4.2.47)$$


which for $z_{\text{cut}} < 1$ gives

$$\sigma_{\text{cone}}^g \sim \pi R^2 \left[\log \left(z_{\text{cut}} \frac{p_T^\gamma}{E_T^{\text{max}}} \right) - \frac{1}{2}\beta + 2 \sum_{k=1}^{\infty} \frac{1}{k} \frac{z_{\text{cut}}^k}{k\beta + 2} \right] \quad (4.2.48)$$


$$\leq \pi R^2 \left[\log \left(z_{\text{cut}} \frac{p_T^\gamma}{E_T^{\text{max}}} \right) - \frac{1}{2}\beta - \frac{2}{\beta + 2} \log(1 - z_{\text{cut}}) \right]. \quad (4.2.49)$$

The correspondence of eq. (4.2.42) between smooth-cone and soft-drop isolation only applies directly when there is only one parton in the photon-cone; when there is more than one, the object being tested against the soft-drop criterion is not a parton (as in the smooth-cone case), but a protojet object comprising more than one parton.

However, in an NNLO calculation there are at most two partons, and the configurations in which both are present in the photon cone (or, more precisely, can be combined into a protojet which falls within the cone), are highly suppressed. In fig. 4.11 we show a selection of differential cross-sections demonstrating that for NNLO calculations and $z_{\text{cut}} = \varepsilon_\gamma = 0.1$, and $\beta = 2n = 2$ the practical difference between smooth-cone and soft-drop isolation is negligible. These parameters correspond to the ‘tight’ isolation most commonly implemented in practice.



(a) Parameter-dependence of soft-drop isolation for quark-photon collinear splittings.



(b) Parameter-dependence of soft-drop isolation with a gluon in the photon cone.

Figure 4.10: The parameter-dependence of soft-drop isolation, according to the identity of the parton isolated against. $\Delta\sigma$ is the difference between the cross-section with the specified isolation parameters and the baseline parameters $R = 0.4, \beta = 2, z_{\text{cut}} = 0.1$ (indicated by the grey lines). The remaining cuts are those of the ATLAS experimental results discussed in chapter 5 and match the corresponding plots figs. 4.3a and 4.3b.




Figure 4.11: Differential distributions at NLO and NNLO showing the agreement between soft-drop and smooth-cone isolation, for cone radius $R = 0.4$, and parameters $\varepsilon_\gamma = 0.1, n = 1$ for smooth-cone isolation (with profile function eq. (4.2.6)), and $z_{\text{cut}} = 0.1, \beta = 2$ for soft-drop isolation.

4.2.4. Infrared sensitivity of cone isolation

In general, any parton-level cone-based isolation criterion of the generic form eq. (4.2.4) amounts to a veto implemented through a measurement function containing factors of the form

$$\prod_{\gamma} \prod_{i=1}^n \mathcal{I}_{\gamma i}, \quad (4.2.50)$$

where the index i ranges over final-state partons, and

$$\mathcal{I}_{\gamma i} = \Theta \left[E_{\text{T}}^{\text{iso}}(\gamma) \chi \left(\min \left(\Delta R_{\gamma i}, R \right); R \right) - \sum_{j=1}^n E_{\text{T}}^j \Theta \left[\min \left(\Delta R_{\gamma i}, R \right) - \Delta R_{\gamma j} \right] \right] \quad (4.2.51)$$

(using the $\Theta(0) = 1$ convention). This is zero, and hence vetoes events, in which the accumulated partonic energy in the cone exceeds the profile function.

It can readily be seen from this formalism that the Heaviside step function implies a discontinuity in the integrand at the bounding surface on which the isolation criteria inequalities are exactly saturated. This is an intrinsic property of veto-based isolation techniques. At NLO, where there is a single parton that can only enter the photon cone of the softer photon, the consequences of this become clearer:

$$\mathcal{I}_{\gamma 1} = \Theta \left[E_{\text{T}}^{\text{iso}}(\gamma) \chi \left(\min \left(\Delta R_{\gamma 1}, R \right); R \right) - E_{\text{T}}^1 \Theta \left[\min \left(\Delta R_{\gamma 1}, R \right) - \Delta R_{\gamma 1} \right] \right] \quad (4.2.52)$$

That is, we expect to have introduced a step-like discontinuity inside the physical region at

$$p_{\text{T}}^{\gamma\gamma} \equiv E_{\text{T}}^1 = E_{\text{T}}^{\text{iso}}(\gamma) \chi \left(\Delta R_{\gamma 1}; R \right) \quad \forall \Delta R_{\gamma 1} \leq R \quad (4.2.53)$$

where the integrand is zero for $p_{\text{T}}^{\gamma\gamma} > E_{\text{T}}^{\text{iso}}(\gamma) \chi \left(\Delta R_{\gamma 1}; R \right)$ and non-zero below it, through the formulation of the isolation criterion. This is precisely the source of the discontinuity visible in fig. 4.8. For the complementary region $\Delta R_{\gamma 1} > R$ where the parton is outside the cone, there is no discontinuity in the integrand: the measurement function eq. (4.2.52) is never zero, and so the isolated and unisolated

integrands are identical everywhere. Conversely, examining instead the region defined by $p_T^{\gamma\gamma} \geq E_T^{\text{iso}}(\gamma)$ we see the same step-like discontinuity arising at $\Delta R_{\gamma 1} = R$, the boundary of the isolation cone.

Such discontinuities within the physical region were first described in general in [209]. For diphoton production they first arise in the NLO-plus-fragmentation calculation and were remarked upon in [192], but have not previously been identified in the NNLO direct production calculation. They represent a localised breakdown of perturbation theory in which a step-like discontinuity leads at higher orders to infrared Sudakov singularities. These arise from the disruption of the expected cancellation between soft real gluons and the corresponding virtual corrections, since isolation vetoes a subset of the former without affecting the latter. Resummation of the generated logarithms is then expected to restore continuity of the distribution, resulting in a characteristic ‘Sudakov shoulder’. Following the logic outlined in [209], the step-like isolation behaviour shown in fig. 4.8 leads to a double-logarithmic divergence in the region $p_T^{\gamma\gamma} < E_T^{\text{thr.}}$,

$$\Delta_{p_T^{\gamma\gamma}=E_T^{\text{thr.}}}^{\pm} \sim -\ln^2 \left[1 - \left(\frac{p_T^{\gamma\gamma}}{E_T^{\text{thr.}}} \right)^2 \right]. \quad (4.2.54)$$

This behaviour does indeed arise in the NNLO $d\sigma/dp_T^{\gamma\gamma}$ distribution as expected. It is shown alongside the corresponding NLO discontinuity in fig. 4.12, together with the corresponding (continuous) smooth-cone distribution. The distinctive double-singularity shape of the hybrid-isolation distribution is as anticipated in [209], and represents a clear deviation from the expected behaviour of the hybrid-isolation distribution on physical grounds from eq. (4.2.10).

There is an additional Sudakov critical point arising from the boundary of the Born kinematic region at $p_T^{\gamma\gamma} = 0$ which would also be expected to require resummation to generate reliable predictions. The practical effect of this additional singularity at small $p_T^{\gamma\gamma}$ is therefore to revise upwards the lower boundary of the region of the $p_T^{\gamma\gamma}$ -distribution at which we might expect NNLO calculations to accurately describe the data. For current experimental binnings, this effect is negligible. The singularity-




Figure 4.12: Discontinuity in the $d\sigma/dp_T^{\gamma\gamma}$ distribution arising from hybrid isolation at NLO with $E_T^{\text{thr.}} = 11$ GeV, and the resulting Sudakov singularity at NNLO.

ies are integrable, and the positive and negative logarithmic contributions typically cancel against each other in a single bin that contains the critical point. However, as the target precision of both experimental data and theoretical predictions increases, these effects may not remain negligible, especially if a bin-edge coincides with the Sudakov critical point.

We briefly remark on the second discontinuity implied by eq. (4.2.52), in the $\Delta R_{\gamma j}$ distribution. The NLO isolation function eq. (4.2.52) implies a discontinuity in $\Delta R_{\gamma 1}$ at the boundary of the isolation cone. At NLO, where each identified jet comprises a single parton, this would lead to a discontinuity in a $\Delta R_{\gamma 2 j_1}$ distribution, were the jet definition set small enough to allow partons to be simultaneously soft enough to be permitted inside the cone by isolation, and hard enough to be identified as a jet. The obvious tension between these two conditions makes this a theoretical, rather than a phenomenological concern. At NNLO, however, the possibility arises for partons soft enough to be permitted inside the cone by the isolation criteria to be combined with harder partons outside the cone, resulting in a jet with $\Delta R_{\gamma j} > R$. The underlying discontinuity at one order and the resulting Sudakov singularities at the next order would then be displaced relative to one another, and would resemble a new phenomenon of unclear origin. These boundary effects can be expected to lead to unphysical results in any fixed-order prediction of photon-jet separation.

At NLO, the nature of the isolation-induced discontinuity shown in fig. 4.8 is specific to hybrid- and fixed-cone isolation with $\varepsilon_\gamma = 0$. The surface defined in eq. (4.2.52) is a surface of constant $p_T^{\gamma\gamma}$, and hence the discontinuity introduced into the integrand remains in the $d\sigma/dp_T^{\gamma\gamma}$ distribution, and at higher orders gives rise to a Sudakov critical point. More generally, for $\varepsilon_\gamma = 0$ a discontinuity in the $p_T^{\gamma\gamma}$ -distribution arises from any interval on which $\chi(r; R)$ is constant.

The discontinuity is fully regulated in smooth-cone isolation in NLO kinematics, since the boundary in $p_T^{\gamma\gamma}$ at which the discontinuity would arise is no longer a constant $E_T^{\text{thr.}}$, but a monotonic function of r , and the threshold of permitted events is spread evenly across $p_T^{\gamma\gamma}$ rather than discretely at a boundary. This masks

the IR critical point and gives a continuous $p_T^{\gamma\gamma}$ distribution (with a discontinuous derivative⁶). However, it instead introduces one into the $p_T^{\gamma\gamma}/\chi(r)$ distribution.

Within hybrid isolation, continuity can be restored to the $p_T^{\gamma\gamma}$ -distribution by, for example, introducing a small non-zero ε_γ . This amounts to a rotation of the boundary surface, and moves the discontinuity from the $p_T^{\gamma\gamma}$ -distribution into the $(p_T^{\gamma\gamma} - \varepsilon_\gamma p_T^{\gamma 2})$ distribution, as shown in fig. 4.13. The resulting Sudakov singularities in the latter distribution at a higher order then manifest themselves in the $p_T^{\gamma\gamma}$ distribution as an unphysical bump resulting from the remainder of the cancellation of positive and negative Sudakov logarithms in each bin.

These discontinuities, and the resulting singularities, are therefore a necessary consequence of cone-based isolation, and can only be moved between distributions, rather than avoided entirely. The effect of the logarithms is not confined to the distribution that is discontinuous at a lower order, but can leak into correlated distributions, where it may be harder to identify.

In general, any observable whose definition is constructed to align with the axis of the step-function will exhibit this threshold behaviour. Where this coincides at a lower order with an observable of physical interest, it is likely to lead to infrared sensitivity. For sufficiently wide histogram bins (including those used for the ATLAS 8 TeV data), the integrable singularities are masked, whilst binnings that combine both critical points, at $p_T^{\gamma\gamma} = 0$ and $p_T^{\gamma\gamma} = E_T^{\text{thr.}}$ into a single bin disguise both

⁶ That differential cross-sections with respect to variables correlated with the argument of a Θ -function will have discontinuous derivatives can be seen by considering the simplified example

$$I = \int_{-1}^1 dx \int_{-1}^1 dy \Theta(x) = 2 \int_{-1}^1 dx \Theta(x). \quad (4.2.55)$$

The ‘differential’ integral $dI/dx = 2\Theta(x)$ is discontinuous at $x = 0$, whilst dI/dy is continuous (here identically 1, but not in general). Transforming to $u = x - \epsilon y$ and performing the integration over y gives

$$\frac{dI}{du} = \begin{cases} 0 & -(1+\epsilon) \leq u \leq -\epsilon \\ 1 + \frac{u}{\epsilon} & -\epsilon \leq u \leq \epsilon \\ 2 & \epsilon \leq u \leq 1-\epsilon \\ 1 + \frac{1-u}{\epsilon} & 1-\epsilon \leq u \leq 1+\epsilon \end{cases} \quad (4.2.56)$$

The discontinuity has been removed from the ‘differential’ integral, but the symptoms of the discontinuity of the integrand persist, in the discontinuous derivative.




Figure 4.13: The differential cross-sections $d\sigma/dp_T^{\gamma\gamma}$ and $d\sigma/d(p_T^{\gamma\gamma} - \varepsilon_\gamma p_T^2)$ at NLO, for hybrid isolation with $E_T^{\text{thr.}} = 8$ GeV and $\varepsilon_\gamma = 0.1$. The discontinuity of the derivative of $d\sigma/dp_T^{\gamma\gamma}$ can be seen to arise from the discontinuity imposed by the isolation Θ -function, and so will be accompanied by Sudakov logarithms at higher orders.

Sudakov critical points entirely, as in fig. 13 of [207].

Given this, it appears that the phenomenological significance of these singularities is limited, provided that deviations from fixed-order predictions in these regions are not misunderstood to have physical significance. This is easier to recognise in distributions such as $p_T^{\gamma\gamma}$ that are directly constrained by photon isolation than it might be where the analogous observable is not of direct physical interest. This is the case, for example, for the photon-plus-jet process, where different experimental cuts and attention to different observables change the relevance of the expected non-analytic behaviour of $p_T^{\gamma j}$.

However, for colourless final-states including the diphoton final state, the differential cross section with respect to the transverse momentum of the identified

final state has particular significance, as it is relied upon by alternative subtraction schemes such as q_T - or N -jettiness subtraction, as outlined in eq. (2.6.17). It is clear from fig. 4.12 that the $p_T^{\gamma\gamma}$ -dependence of the cross-section at small $p_T^{\gamma\gamma}$ is sensitive to the details of the isolation used and not universal. If the dependence is not universal due to substantial power corrections, the calculation retains a dependence upon the technical cut $r_{\text{cut}} = q_T^{\gamma\gamma}/M_{\gamma\gamma}$ rather than plateauing at the correct value. This would explain the absence of a plateau in the r_{cut} -dependence plots for diphoton production using q_T -subtraction with MATRIX in [210].

These power corrections have been explored analytically in [123], where it was found that for smooth-cone isolation, they grow in magnitude as $(Q/E_T^{\text{iso}})^{1/n}$, with a proposal for how they could be accounted for. As a result, the phenomenological significance of these power corrections should grow as we move to higher centre-of-mass energies, or for tightening isolation. It remains to be seen whether they will pose a meaningful problem for these alternative subtraction schemes.

4.3. Cut-dependence

Finally, we turn to the cut-dependence of the cross-section before proceeding to phenomenology. Since

$$\sigma = \int_{p_T^{\gamma_2, \text{cut}}}^{\infty} \frac{d\sigma}{dp_T^{\gamma_2}} dp_T^{\gamma_2}, \quad (4.3.1)$$

it follows that

$$\frac{\partial \sigma}{\partial p_T^{\gamma_2, \text{cut}}} = - \left. \frac{d\sigma}{dp_T^{\gamma_2}} \right|_{p_T^{\gamma_2} = p_T^{\gamma_2, \text{cut}}}, \quad (4.3.2)$$

and likewise for $p_T^{\gamma_1, \text{cut}}$.

The cross-section should therefore decrease as $p_T^{\gamma_1, \text{cut}}$ or $p_T^{\gamma_2, \text{cut}}$ increases, as the cuts exclude a larger volume of phase-space—except in the Born kinematics, where $d\sigma/dp_T^{\gamma_2} \equiv 0$ for $p_T^{\gamma_2} < p_T^{\gamma_1, \text{cut}}$, so the cross-section is independent of $p_T^{\gamma_2, \text{cut}}$ (as transverse momentum conservation requires $p_T^{\gamma_2} = p_T^{\gamma_1}$).

The corresponding plot of σ against the cuts is shown in fig. 4.14. It shows the expected behaviour for $p_T^{\gamma_2} < p_T^{\gamma_1}$, with the NNLO K -factor varying from a factor of 3 for symmetrical p_T cuts $p_T^{\gamma_1, \text{cut}} = p_T^{\gamma_2, \text{cut}}$, to a factor of 9 for the most asymmetrical cuts plotted,


$$p_T^{\gamma_1, \text{cut}} = 40 \text{ GeV} \qquad p_T^{\gamma_2, \text{cut}} = 20 \text{ GeV}. \quad (4.3.3)$$

Notably, the expected monotonicity is violated in the neighbourhood of symmetrical cuts, with a distinctive unphysical cusp.


This cusp is a general phenomenon that was first discussed in [208] in the context of dijet photoproduction. It was discussed for diphoton production at 7 TeV in [207], where it was found that, for cuts $p_T^{\gamma_2, \text{cut}}$ and $p_T^{\gamma_1, \text{cut}} = p_T^{\gamma_2, \text{cut}} + \delta$, the NLO cross-section behaves as

$$\sigma^{\text{NLO}}(\delta) - \sigma^{\text{NLO}}(0) \propto \delta \log^2 \delta + \mathcal{O}(\delta^2), \qquad \delta > 0 \quad (4.3.4)$$

with this double-logarithmic behaviour arising from the emission of gluons that are



(a) Dependence of the diphoton cross-section σ on the photon p_T cuts.



(b) Dependence of the diphoton K -factor on the photon p_T cuts. The $Z(\rightarrow 2\ell)$ -production K -factors are also shown for reference, for matching final-state cuts, with the p_T -cuts applied to the leptons.

Figure 4.14: The effect of the p_T -cuts on perturbative convergence. In each case, one of the $p_T^{\gamma/\ell}$ cuts is 40 GeV whilst the other is as displayed on the x -axis. The bands indicate 7-point scale variation about the central scale $M_{\gamma\gamma}$ ($M_{\ell\ell}$).

both collinear with the initial-state partons, and soft. This behaviour implies that at the cusp itself the derivative with respect to δ diverges. This can be seen in fig. 4.14.


In fig. 4.15 we show the corresponding behaviour for the $q\bar{q}$ and qg channels. As expected, the ‘divergent-derivative’ cusp is only present at NLO in the $q\bar{q}$ channel, where the emission of final-state soft-collinear gluons is possible. At NNLO, the cusp is softened but the unphysical behaviour persists, since each additional parton introduces additional double logarithms of the same origin (suppressed by additional factors of α_s). The removal of this non-analytic and unphysical behaviour therefore requires resummation [211]. The qg -channel also shows unphysical behaviour, violating monotonicity, but with only a finite discontinuity of the derivative at the point of symmetrical cuts.

Through eq. (4.3.2), we can connect the unphysical behaviour at the cusp to the unphysical behaviour of $d\sigma/dp_T^{\gamma_2}$ in fig. 4.8. This region is sensitive to isolation, as $p_T^{\gamma_2}$ lies within the strip of eq. (4.2.33), $p_T^{\gamma_2} \in [29, 40]$ GeV. This implies that the derivative of σ as a function of the cuts should be discontinuous at $p_T^{\gamma_{1,\text{cut}}} = 51$ GeV, and at $p_T^{\gamma_{2,\text{cut}}} = 29$ GeV, arising from the Θ -function imposing the photon isolation.


Figures 4.15a and 4.15b additionally shows that the channel-breakdown of the calculation is highly sensitive to the asymmetry of the cuts, through which the qg -channel can be enhanced or suppressed relative to the $q\bar{q}$ -channel. In particular, we see that for the cuts used by ATLAS,

$$p_T^{\gamma_{1,\text{cut}}} = 40 \text{ GeV} \qquad p_T^{\gamma_{2,\text{cut}}} = 30 \text{ GeV}, \qquad (4.3.5)$$

to which we shall turn in chapter 5, the large K -factor (of almost 6) shown in fig. 4.14b is chiefly attributable to the effect of the asymmetry of the chosen cuts on the qg - and gq -channel cross-sections.



(a) Dependence of $\sigma_{q\bar{q}}$ on the photon p_T cuts. (b) Dependence of σ_{gg} on the photon p_T cuts.



(c) Dependence of the $q\bar{q}$ -channel K -factor on the photon p_T cuts. (d) Dependence of the gg -channel K -factor on the photon p_T cuts [ratio to $\sigma_{q\bar{q}}^{\text{LO}}$ for comparison].

Figure 4.15: The effect of the p_T -cuts on perturbative convergence in the $q\bar{q}$ - and gg -channels. Note that here, the (identical) contributions from the $q\bar{q}$ - and gg -channels are excluded; their inclusion would consistently multiply all the results in figs (a) and (b) by a factor of 2, and leave the K -factors unchanged.

Diphoton production: phenomenology

In this chapter we will apply the NNLOJET calculation outlined in section 3.5 to collider phenomenology, calculating cross-sections and differential cross-sections corresponding to experimental measurements. Understanding the applicability of the calculation to the experimental measurements requires the theoretical understanding developed in chapter 4.

We shall focus on the ATLAS 8 TeV data [212] from Run I of the LHC, based on an integrated luminosity of 20.2 fb^{-1} and released in 2016. The relevant fiducial cuts are:

$$p_T^{\gamma_1} > 40 \text{ GeV} \qquad p_T^{\gamma_2} > 30 \text{ GeV} \qquad (5.0.1a)$$

$$\Delta R_{\gamma\gamma} > 0.4 \qquad |y^\gamma| \in [0, 1.37) \cup (1.56, 2.37) \qquad (5.0.1b)$$

$$E_T^{\text{iso,part}} < 11 \text{ GeV} \qquad \text{within cone } \Delta R \leq 0.4. \qquad (5.0.1c)$$

In figs. 5.1 and 5.2 we reproduce several of the plots that accompanied the data's release, and which show the NNLO prediction of 2γ NNLO substantially underestimating the measured cross-section.

We choose $R_d = 0.1$ for hybrid isolation, as outlined in section 4.2.2.2, and smooth-cone isolation parameters $n = 1$ and $E_T^{\text{iso}} = 11 \text{ GeV}$ for both χ and χ^{hyb} . Here and throughout we use the NNPDF 3.1 parton distribution functions [53]. The




Figure 5.1: Measured ATLAS fiducial cross section, compared to predictions from SHERPA 2.2.1, DIPHOX, RESBOS and $2\gamma\text{NNLO}$. Taken from [212].




Figure 5.2: Differential cross-sections $d\sigma/dM_{\gamma\gamma}$ and $d\sigma/d|\cos \theta_\eta^*|$ as measured by ATLAS, compared to an NNLO prediction computed with $2\gamma\text{NNLO}$ and an NLO + parton shower calculation computed by SHERPA. Taken from [212].

QED coupling constant α is set at $\alpha_{\text{em}}(0) = 1/137$.

5.1. Hybrid isolation for phenomenology

5.1.1. R_d dependence

The $\log R/R_d$ scaling of eq. (4.2.30) indicates that the cross-section diverges in the small-inner-cone limit, as can be seen in fig. 4.5, and as expected from the discussion of section 4.2.1.1. This arises because the partition of phase-space into a cone of radius R and its complement induces $\log R$ contributions in both, which cancel in their sum. Any isolation procedure applied only inside the photon cone changes the former but not the latter, leading to a miscancellation of logarithms, the remainder of which will become large in the small- R limit.

We must therefore be careful to choose a value of R_d that is large enough to regulate the collinear singularity, but small enough to approximate the fixed-cone result better than the smooth-cone value $R_d = R$. Ideally, this would be approximately equal to that at which the compensation that was discussed above occurs, to reproduce the cross-section given by fixed-cone isolation.

To determine the value of R_d at which this compensation occurs, in [2] we compared NLO cross-sections and differential distributions obtained at fixed order with hybrid isolation to those obtained using DIPHOX [192] with fixed-cone isolation. fig. 5.3 shows that for the ATLAS-motivated cuts $E_T^{\text{thr.}} = 11 \text{ GeV}$, $R = 0.4$ and $R_d = 0.1$ the hybrid isolation result was almost fully contained within the DIPHOX uncertainty band, except where the fragmentation contribution populated regions of phase-space that first enter the fixed-order calculation at the subsequent order of perturbation theory.




Figure 5.3: Comparison between direct diphoton production at NLO with hybrid isolation, against the fragmentation-inclusive calculation with standard isolation computed by DIPHOX. The parameters $n = 1$ and $R_d = 0.1$ give good agreement. From [2].

5.1.2. Comparison of hybrid and smooth-cone distributions

We first explore the effect of moving from smooth-cone to hybrid isolation on differential cross-sections chosen to illustrate the underlying features.

In fig. 5.4 we show $d\sigma/d\Delta R_{\gamma\gamma}$ and $d\sigma/dM_{\gamma\gamma}$. The relative enhancement is greatest at low $M_{\gamma\gamma}$, whilst the absolute enhancement $d\Delta\sigma/dM_{\gamma\gamma}$ follows the shape of the underlying distribution, with the difference largest at the Born threshold of $M_{\gamma\gamma} = 80 \text{ GeV}$. Broadly these reflect the two dominant configurations in which soft partonic emissions can enter into isolation cones: either the photons are balanced against each other (Born-like), or the diphoton system is relatively collimated and balanced against a jet. Accordingly, configurations with the explicit requirement of an extra jet see a further peak in $d\Delta\sigma/dM_{\gamma\gamma}$ at $M_{\gamma\gamma} \approx p_T^{j_{\text{cut}}}$ corresponding to $\Delta R_{\gamma\gamma} \approx \Delta R_{\gamma\gamma}^{\text{cut}}$, as shown in fig. 5.5. In the $p_T^{j_{\text{cut}}} \rightarrow 0$ limit, this is effectively truncated by the cuts on $\Delta R_{\gamma\gamma}$, which is the configuration corresponding to the



(a) $\Delta R_{\gamma\gamma}$ at NLO and NNLO using matched-hybrid and smooth-cone isolation.(b) The induced effects at low $\Delta R_{\gamma\gamma}$ on $M_{\gamma\gamma}$.

Figure 5.4: $d\sigma/\Delta R_{\gamma\gamma}$ and $d\sigma/M_{\gamma\gamma}$ at NLO and NNLO using matched-hybrid and smooth-cone isolation. The deviations for small $M_{\gamma\gamma}$ and $\Delta R_{\gamma\gamma}$ are related, as events with small $M_{\gamma\gamma}$ can only both pass the photon cuts if they have sufficiently small $\Delta R_{\gamma\gamma}$. For example, for these cuts, $M_{\gamma\gamma} \leq 27 \text{ GeV}$ requires $\Delta R_{\gamma\gamma} \leq 0.8$.

small $\Delta R_{\gamma\gamma}$ effects seen in fig. 5.4. As the jet cut increases this peak will become dominant.

The requirement of a jet imposes a lower bound on $p_T^{\gamma\gamma}$ and so removes the Sudakov instabilities of the inclusive distribution that were discussed in section 4.2.4. The two peaks in the two plots correspond to the same physics in the opposite order, with the peak at $p_T^{\gamma\gamma} \approx p_T^{j\text{cut}}$ corresponding to the configuration in which the photons and jet are balanced, and the second peak at $p_T^{\gamma\gamma} \approx 75 \text{ GeV}$ corresponding to the threshold at 70 GeV, the smallest value that can be generated within the cuts for every value of $\Delta\phi_{\gamma\gamma}$, shown in fig. 5.6. Below this threshold the photon cuts imply an implicit minimum for $\Delta\phi_{\gamma\gamma}$, restricting the available phase-space. Contributions from this peak give rise to a distinctive cusp in both the experimental and the NNLO distributions which, corresponding to the small- $\Delta R_{\gamma\gamma}$ region, is especially sensitive to isolation. Smooth-cone isolation suppresses the kinematic peak in this region, which is restored by the less restrictive hybrid isolation profile function.

Finally, for completeness, in fig. 5.7 we consider four further differential cross-sections of interest. The $p_T^{\gamma_1}$ and $p_T^{\gamma_2}$ distributions are affected most substantially at the boundary of the photon cuts, as expected from fig. 4.8, but are elsewhere mostly unchanged by modifications of the cuts. These regions dominate the cross-section, and explain the large R_d -sensitivity of fig. 4.9. Whether the correction here is purely physical or, particularly for the $p_T^{\gamma_2}$ distribution, arises from unphysical behaviour at the boundary of the Born phase-space, is unclear. As at NLO, for the rapidity separation $\Delta y_{\gamma\gamma}$ the additional events permitted by hybrid isolation amount to an overall constant factor in $d\Delta\sigma/d\Delta y_{\gamma\gamma}$.

In this section we have compared smooth-cone to hybrid isolation at NNLO for a range of differential cross-sections of phenomenological significance. The effect of interchanging them, which indicates the uncertainty associated to the theoretical implementation of the isolation criteria, is substantial and leads to effects of approximately 10% in uncorrelated distributions, and localised effects of up to 40% in distributions highly sensitive to the specifics of the isolation criteria. Uncertainties




Figure 5.5: The absolute difference between the hybrid- and smooth-cone isolation differential cross-sections $d\sigma/dM_{\gamma\gamma}$ and $d\sigma/dp_T^{\gamma\gamma}$ for $R_d = 0.1, 0.2, 0.3$, for diphoton production in association with an anti- k_T jet of $p_T \geq 25$ GeV, with $R = 0.4$.




Figure 5.6: The dependence of $p_T^{\gamma\gamma}$ on $\Delta\phi_{\gamma\gamma}$, for fixed $p_T^{\gamma_1} = 40$ GeV and $p_T^{\gamma_2} = 30$ GeV on the threshold of the ATLAS cuts. Accordingly, for values of $p_T^{\gamma\gamma}$ below 70 GeV there is an implicit constraint on $\Delta\phi_{\gamma\gamma}$, here illustrated for $p_T^{\gamma\gamma} = 45$ GeV, arising from the photon p_T -cuts. The full angular range of $\Delta\phi_{\gamma\gamma}$ can only contribute for $p_T^{\gamma\gamma} \geq 70$ GeV.




Figure 5.7: The NNLO distributions $d\sigma/dp_T^{\gamma_1}$, $d\sigma/dp_T^{\gamma_2}$, $d\sigma/dp_T^{j_1}$ and $d\sigma/d\Delta y_{\gamma\gamma}$ for hybrid-cone and smooth-cone isolation respectively, and the ratio between the smooth-cone and the hybrid distributions. The defining jet requirement for the third plot is of an anti- k_T jet with $p_T^j \geq 25$ GeV and $R = 0.4$.

of this magnitude are compatible with the size of the scale-uncertainty band, and therefore represent a substantial theory uncertainty that should be accounted for.

We will return to consider isolation effects in tandem with scale choice in section 5.3.

5.2. Scale choice

A further uncertainty in the theoretical calculation arises from the choice of functional form μ_0 for the renormalisation and factorisation scales. The conventional choice is $\mu_0 = M_{\gamma\gamma}$, the invariant mass of the diphoton system, with the magnitude of missing higher-order-uncertainties (MHOUs) estimated through the envelope of the variation $\mu_{R,F} = \xi_{R,F} \cdot \mu_0$ for $\xi_R, \xi_F \in \{\frac{1}{2}, 1, 2\}$.

Where two *a priori* reasonable choices of μ_0 themselves differ by a factor greater than 2, either locally or globally, this procedure fails to span the uncertainty of the calculation even at the known orders. Any estimate of MHOUs is therefore potentially unreliable.

We begin by briefly reviewing the common scale choices for related processes. In sections 5.2.2 and 5.2.3 we then look at the effects of moving between two choices motivated by these, $\mu_0 = M_{\gamma\gamma}$ and $\mu_0 = \langle p_T^\gamma \rangle$, the arithmetic mean of the photon transverse momenta of the two required photons. Finally, in section 5.2.4 we generalise to a wider class of possible scale choices.

5.2.1. Scale choice for photon processes

We briefly summarise the scale choices used in the literature for this and related processes. In [192], the first NLO study of diphoton production with fragmentation (DIPHOX), the authors used $\mu_0 = \frac{11}{20} \langle p_T^\gamma \rangle$ for fixed-target data, and $\mu_0 = M_{\gamma\gamma}$ as the central scale for LHC predictions. This scale is also used for NNLO calculations making predictions for or comparisons with data in [202, 203, 207] and the experimental papers applying them to measurements at the Tevatron [213, 214] and the

LHC [197, 212, 215]. In [207] the scale $\mu_0 = M_T^{\gamma\gamma} = \sqrt{M_{\gamma\gamma}^2 + (p_T^{\gamma\gamma})^2}$ is additionally considered, finding that the results differ from those for $M_{\gamma\gamma}$ only in regions of distributions that correspond to the presence of a hard, high- p_T jet.


For an inclusive single photon and a single photon in association with a jet, p_T^γ is used in the NNLO calculations of [165, 216]. In the context of PDF fits, it was found in [217] that direct photon production data with the former NNLO calculation and scale p_T^γ could be incorporated into the NNPDF 3.1 global fit without exhibiting tensions with other data.

For triphoton production, $M_{\gamma\gamma\gamma}$ is used for the MCFM NLO calculation in [158], and $\frac{1}{4}H_T = \frac{3}{4}\langle p_T^\gamma \rangle$ and $\frac{1}{2}H_T = \frac{3}{2}\langle p_T^\gamma \rangle$ are both found to be in agreement with data in the NNLO calculation of [204].

Finally, we note that the closest kinematically-related process whose measurements were used in the NNPDF 3.1 fit is that of single-inclusive jets, for which the jet p_T was used as the central scale. A more recent study of the scale-choice for single-inclusive jet cross-sections [218] used the central choice $\hat{H}_T$, the scalar sum of the transverse momenta of all partons in the event.

This illustrates that the conventional choice for diphoton production of $\mu_0 = M_{\gamma\gamma}$ is somewhat atypical among related processes. Its main advantage is for Higgs processes or through the analogy with dilepton final states arising from heavy-boson decay. For such processes the invariant mass of the conditioned-upon two-particle final-state particles gives the imputed invariant mass of the virtual boson. For QCD photon production, however, there is no particle to which this invariant mass is expected to correspond, and no QCD vertex with which it can be associated. To explore the significance of this convention we therefore choose to compare $\mu_0 = M_{\gamma\gamma}$ against alternatives below, focusing on $\mu_0 = \langle p_T^\gamma \rangle$.¹

¹ We have already seen in section 2.4.3 that choosing the geometric mean of several scales reduces a multi-scale problem to a single-scale problem. The arithmetic mean does not have this property (though it will do so approximately for the dominant region of phase-space), but has the advantage of being intuitively easy to understand.



(a) The kinematic region $(x, M_{\gamma\gamma}^2)$ probed within the ATLAS cuts. (b) The kinematic region $(x, \langle p_T^\gamma \rangle^2)$ probed within the ATLAS cuts.

Figure 5.8: The kinematic regions (x, Q^2) probed by diphoton production at leading-order (Born) and higher-orders (HO), according to the choice of scale. These are the arguments of the PDFs in eq. (2.1.3) with the corresponding choices of factorisation scale μ_F .

5.2.2. Perturbative convergence

We first consider the perturbative convergence of the cross-section. In fig. 5.9 we show the cross-section and K -factors at NLO and NNLO for a number of choices of dynamic scale, as well as the scale evolution calculated from the renormalisation group equations.

The NLO K -factors are consistently large due to the opening of the qg channel and the asymmetry of the cuts (as explained in section 4.3), and vary according to its considerable dependence on the scale choice. The K -factor for the $q\bar{q}$ channel alone is approximately 1.5. The cross-sections for dynamic scale choices are largely consistent with the fixed-scale calculation corresponding to their mean value, suggesting the reweighting of phase-space by the dynamism of the dynamic central scales has a limited effect on the total cross-section. At NNLO, the K -factor is still considerable (approximately 1.4), due to sizeable NLO corrections in the qg channel (K -factor ~ 1.3), NNLO corrections in the $q\bar{q}$ channel (~ 1.2), and the opening of the gg and

qq' channels, but is stable for all the choices of scales considered.

Overall, as expected from the running of α_s , dynamic scales which range over smaller values lead to larger predictions than those with larger values. Purely in terms of the distribution of their magnitude, the scales $\langle p_T^\gamma \rangle$ and $M_{\gamma\gamma}$ represent the two extremes between which other reasonable dynamic scales are likely to fall.

Despite the stability of the NNLO-to-NLO K -factor across these choices of scales, it is clear from the gradient of the grey band that the scale-dependence remains significant. The use of a dynamic rather than a fixed scale can be seen to bring the scales into closer agreement than would be expected from their central values alone.

5.2.3. Kinematic effects

We now consider the kinematics of the two scales $M_{\gamma\gamma}$ and $\langle p_T^\gamma \rangle$, focusing on regions of phase-space in which we expect the ratio $M_{\gamma\gamma}/\langle p_T^\gamma \rangle$ to become large (or small) and potentially lead to discrepancies arising from large logarithms of ratios of the scales. Although we focus on the diphoton context, including the ATLAS cuts, the underlying kinematic properties are universal.

In the Born kinematics, $\langle p_T^\gamma \rangle = p_T^{\gamma_1} = p_T^{\gamma_2}$ and

$$M_{\gamma\gamma} = 2 \langle p_T^\gamma \rangle \cosh \left(\frac{1}{2} \Delta y_{\gamma\gamma} \right) \geq 2 \langle p_T^\gamma \rangle. \quad (5.2.1)$$

The fiducial cuts on rapidity separation restrict $|\Delta y_{\gamma\gamma}| \leq 4.74$ and hence in the Born kinematics,

$$2 \langle p_T^\gamma \rangle \leq M_{\gamma\gamma} \leq 10.8 \langle p_T^\gamma \rangle. \quad (5.2.2)$$

Thus already at leading order, the two scales differ by at least the factor of 2 used in the conventional renormalisation and factorisation scale variation. We can therefore anticipate there to be regions of differential distributions in which the scale uncertainty bands around the two choices of μ_0 do not overlap.

The exponential behaviour of the scale $M_{\gamma\gamma}$ at high rapidity separations persists




Figure 5.9: Scale dependence of the cross-section. Each cross-section associated with a dynamical scale μ_0 is plotted against its mean, $\langle \mu_0 \rangle$. The x -error bars indicate the lower- and upper-quartiles of the scale-variable distribution, calculated from the binned data. The scale uncertainties associated with the usual 7-point scale variation around the central scale. The grey bands give the cross-section for the fixed scales μ_R, μ_F specified, calculated from the renormalisation group equations eq. (2.4.3).

to all orders, with the general expression, plotted in fig. 5.10,

$$M_{\gamma\gamma} = \sqrt{2p_T^{\gamma_1} p_T^{\gamma_2} (\cosh \Delta y_{\gamma\gamma} - \cos \Delta \phi_{\gamma\gamma})}. \quad (5.2.3)$$

At higher orders, $M_{\gamma\gamma} \leq \langle p_T^\gamma \rangle$ becomes possible. $M_{\gamma\gamma}$ is bounded below only as a result of the photon separation cut $\Delta R_{\gamma\gamma} \geq 0.4$, which restricts

$$M_{\gamma\gamma} \geq 2\sqrt{p_T^{\gamma_1} p_T^{\gamma_2}} \sin\left(\frac{1}{2}\Delta R_{\gamma\gamma}^{\text{cut}}\right) > 13.76 \text{ GeV} \quad (5.2.4)$$

for the ATLAS cuts described in eq. (5.0.1). Without this cut, which is set to be equal to the isolation cone radius by experiment specifically to exclude each photon from the isolation cone of the other, $M_{\gamma\gamma}$ would in principle be permitted within the calculation to get arbitrarily small. Thus for fixed $p_T^{\gamma_1}$ and $p_T^{\gamma_2}$ (and hence fixed $\langle p_T^\gamma \rangle$), $M_{\gamma\gamma}$ can vary over a factor of approximately 25:

$$0.397 \leq \frac{M_{\gamma\gamma}}{\sqrt{p_T^{\gamma_1} p_T^{\gamma_2}}} \leq 10.8 \quad (5.2.5)$$

with the size of this factor entirely dependent on cuts chosen for primarily experimental reasons. Were the photon-separation cut allowed to become smaller (e.g. to $\Delta R_{\gamma\gamma} \geq 0.2$), or the maximum rapidity separation allowed to grow (e.g. from 4.74 to 6), this ratio would span two orders of magnitude.

To illustrate the range of values taken by the ratio $M_{\gamma\gamma}/\langle p_T^\gamma \rangle$ we show the corresponding normalised distribution at LO, NLO and NNLO in fig. 5.11. We see that the modal value for the ratio is 2, and that the regions where the logarithm of the ratio will be large are suppressed in their contribution to the cross-section, and predominantly arise from the NNLO contribution as additional partonic radiation allows the kinematic configuration to depart further from the Born.

The distortive effect of the scale choice on differential cross-sections depends substantially on the order of the strong coupling α_s , through the renormalisation group equations. This is illustrated in fig. 5.12. The $d\sigma/d\Delta y_{\gamma\gamma}$ distribution exposes the exponential behaviour remarked upon in eq. (5.2.3). At leading-order α_s^0 , the calculation is independent of μ_R , and so the dependence is only on μ_F through




Figure 5.10: The exponential dependence of $M_{\gamma\gamma}$ on the rapidity separation of the photons, for fixed $p_T^{\gamma_1} = 40$ GeV and $p_T^{\gamma_2} = 30$ GeV on the threshold of the ATLAS cuts. The rapidity range of the x -axis is precisely that permitted by the ATLAS cuts, but the combined limit $\Delta\phi_{\gamma\gamma}, \Delta y_{\gamma\gamma} \rightarrow 0$ is prohibited by the cut $\Delta R_{\gamma\gamma} \geq 0.4$.

the PDFs. The dependence is mild: the results for the scale choice $\mu_0 = M_{\gamma\gamma}$ are modestly larger than those for $\mu_0 = \langle p_T^{\gamma} \rangle$, with the deviation largest for $\Delta y_{\gamma\gamma} = 0$ where $M_{\gamma\gamma} = 2 \langle p_T^{\gamma} \rangle$ exactly, due to eq. (5.2.1), and as can be seen through the coincidence of the scale bands of one scale with the central scale of the other.

Additional powers of the coupling constant α_s reverse that hierarchy, due to the monotonicity of the running of the coupling that ensures $\alpha_s(\mu_1) \geq \alpha_s(\mu_2)$ for $\mu_1 \leq \mu_2$. Thus in the regions of large rapidity-separation, the $\mu_0 = M_{\gamma\gamma}$ predictions are suppressed relative to those for $\mu_0 = \langle p_T^{\gamma} \rangle$ by up to 30%.

In the extremes of the distribution, this is driven by the constructive interference of factorisation- and renormalisation-scale variation, in the sense that larger μ_F and larger μ_R both act to suppress the result. The substantial correlation between $\Delta y_{\gamma\gamma}$ and $M_{\gamma\gamma}$ in these bins leads to an implicit cut on $M_{\gamma\gamma}$ in each bin, which leads to artificially small scale-uncertainty bands for the $\mu_0 = M_{\gamma\gamma}$ result compared to variation over an inclusive dynamical scale variable. This might lead to the conclusion




Figure 5.11: Contributions to the cross-section according to the ratio of scale variables $M_{\gamma\gamma}$ and $\langle p_T^\gamma \rangle$. As phase-space constraints are lifted by the emission of additional partons, large ratios of scales become possible.

that the $\mu_0 = M_{\gamma\gamma}$ distributions display improved perturbative convergence due to the narrower scale bands, when it is in fact an artefact of correlation of the scale with the binned observable, leading to a restricted domain for the scale variation procedure.

The behaviour of the $d\sigma/d\Delta R_{\gamma\gamma}$ distribution at low $\Delta R_{\gamma\gamma}$ shows exactly the inverse behaviour: small values of the ratio $M_{\gamma\gamma}/\langle p_T^\gamma \rangle$ lead to an enhanced distribution. As discussed in section 5.1.2, the low- $M_{\gamma\gamma}$ distribution corresponds exactly to small values of $\Delta R_{\gamma\gamma}$, as a result of the cut on photon transverse momenta. This accounts for the common behaviour between the bottom two plots. For an event in the lowest $M_{\gamma\gamma}$ -bin, the NNLO contributions to the cross-section with the scale $\mu_0 = M_{\gamma\gamma}$ are weighted relative to the $\mu_0 = \langle p_T^\gamma \rangle$ contribution with a factor proportional to the ratio of α_s^2 evaluated at the two scales, which is imperfectly compensated by the corresponding dependence in the real-virtual matrix elements. This gives rise to the extreme $\sim 30\%$ deviations between the scale choices in this region; the factorisation-scale dependence is negligible. Since the lower bound on $M_{\gamma\gamma}$ is set by the experimental $\Delta R_{\gamma\gamma}$ and p_T^γ cuts rather than any theory considerations, smaller values of these cuts would lead to still greater distortions between the scale choices. Note that this is in contrast to the problem of scale choices for the dijet process, in which scale choices M_{jj} and $\langle p_T^j \rangle$ differ substantially at NLO but less so at NNLO [218].

5.2.4. Alternative scale functional forms

We remark on the elements of the above discussion which carry over to scale choices with functional forms other than $\mu_0 = M_{\gamma\gamma}$ and $\mu_0 = \langle p_T^\gamma \rangle$. Popular candidates commonly found in studies of other processes typically involve a weighted average, mixing four-momentum-invariant-type observables with transverse-plane observables, schematically of the form

$$\mu_0 = \left(\alpha M_{\gamma\gamma}^r + \beta f \left(\left\{ \mathbf{p}_{T,i} \right\} \right)^r \right)^{\frac{1}{r}} \quad (5.2.6)$$




Figure 5.12: Order-by-order comparison of the difference between the scale choice $\mu_0 = M_{\gamma\gamma}$ and $\mu_0 = \langle p_T^\gamma \rangle$. Leading order here means α_s^0 , so the counterbalancing effects of the PDFs and the running of α_s can be deduced.

where common choices for f include $p_T^{\gamma\gamma}$, the transverse momentum of the diphoton system, or the total transverse momentum of all partons, all jets, or both photons. A variety of functional forms of this type were considered in [219] for the production of a photon pair in association with up to three identified jets.

Functional forms containing $M_{\gamma\gamma}$, i.e. with $\alpha \neq 0$, are dominated by the exponential function of rapidity separation in the (sufficiently) large rapidity-separation region discussed above, and so behave like $M_{\gamma\gamma}$ there. The results in this region therefore lie within the envelope bounded by the scale variation $\mu \in \left\{\frac{1}{2}M_{\gamma\gamma}, 2M_{\gamma\gamma}\right\}$. Of particular importance is the choice

$$\begin{aligned}\mu_0 = M_{T,\gamma\gamma} &= \sqrt{M_{\gamma\gamma}^2 + (p_T^{\gamma\gamma})^2} \\ &= \sqrt{(p_T^{\gamma_1})^2 + (p_T^{\gamma_2})^2 + 2p_T^{\gamma_1}p_T^{\gamma_2} \cosh \Delta y_{\gamma\gamma}}\end{aligned}\quad (5.2.7)$$

which was considered in [207] and shows identical behaviour in this limit. In addition, the scales H'_T , $\hat{H}'_T$, $\sqrt{\Sigma^2}$ and $\sqrt{\hat{\Sigma}^2}$ that were investigated for diphoton production in association with up to three jets in [219] all have similar behaviour, as a consequence of their dependence on $M_{\gamma\gamma}$.

From a physical perspective, the behaviour in this limit represents the scale ambiguity between transverse-plane and four-momentum observables. For central final-states, both classes of observable are of the same order of magnitude and induce a similar ordering of events by scale. For events with large rapidity separation, the projection onto the transverse plane dramatically changes the apparent energy scale of the event. In the extremes of rapidity separation we enter the two-large-scales regime, in which resummation or other approaches may become relevant to correct for large logarithms of the form $\ln(\hat{s}/E_T^2)$. It is possible that compensating behaviour partially accounting for these logarithms would arise in the parton distribution functions if one or the other type of scale was used consistently in fits.

The second region discussed above, of small $\Delta R_{\gamma\gamma}$, arises as a direct consequence of the specific form of the angular factor $(\cosh \Delta y_{\gamma\gamma} - \cos \Delta \phi_{\gamma\gamma})$ in $M_{\gamma\gamma}$, which reduces to $\Delta R_{\gamma\gamma}$ in this limit. As a result, modifying $M_{\gamma\gamma}$ by any offset function

f with a non-zero limit as $\Delta R_{\gamma\gamma} \rightarrow 0$ rectifies the problematic behaviour. This is the case for $M_{T,\gamma\gamma}$ in eq. (5.2.7) above, and all other scales considered with non-zero β . Whilst candidates for f with similar asymptotic behaviour to $M_{\gamma\gamma}$ do exist (e.g. $f = p_T^{\gamma_1} p_T^{\gamma_2} \Delta R_{\gamma\gamma}$), they do not arise naturally from a consideration of the scale of the process. From a physical perspective, problematic behaviour in this region can be explained as the failure of the scale $M_{\gamma\gamma}$ to capture the natural scale of the underlying process, in which the collimated diphoton pair recoils against a hard jet. The scale variable vanishes as the two photons become collinear, restricted only by the experimental cut, even as the event maps onto a photon-plus-jet event of characteristic scale $p_T^\gamma \sim p_T^j$. This leads to exaggerated contributions from α_s which are not compensated by the real-virtual matrix elements.

We can understand this substantial exposure as follows. The diphoton final-state is a two-particle final-state, so the Born-level kinematics are highly restricted; it is colourless, so only the $q\bar{q}$ -channel is fully NNLO, and there is no resonant propagator, so the cross-section is not dominated by a single modal value of the final-state invariant mass. It might therefore be expected that other final-states are unlikely to yield similar sensitivities. Nevertheless, with the same cuts, the same ratios of scales would arise for, e.g., the $Z \rightarrow 2\ell$ process, and it may be worth investigating their impact further.

5.3. Combined effect of isolation and scale variation

Finally we illustrate the combined effect of the simultaneous variation of scale and isolation choice on the distributions. We have previously seen in fig. 5.4 that the region of phase-space most affected by the difference between smooth-cone and hybrid isolation is that in which $\Delta R_{\gamma\gamma}$ is small, and that the same region is highly sensitive to the scale choice, growing starkly with the running coupling relative to a prediction using a scale independent of $\Delta R_{\gamma\gamma}$.




Figure 5.13: Combined ratio plots for four-way scales and isolation comparison, at NNLO (ratio to $\mu_F = \mu_R = \langle p_T^\gamma \rangle$ with hybrid isolation).

We therefore examine the relative size of these competing effects in fig. 5.13. In the top panel, suppression of the cross-section for smooth-cone isolation as $\Delta R_{\gamma\gamma} \rightarrow 0$ competes with the enhancement from the scale $M_{\gamma\gamma}$ to leave the ratio almost flat. As a result, for this specific combination of isolation procedure and scale choice, the competing effects of each choice shown in the lower two panels are disguised, leaving distributions that differ by an overall normalisation.

Away from this region, which is the region not populated by the Born kinematics, the ratio is stable.

5.3.1. Comparison to ATLAS data: four-way comparison

In this section we compare the four combinations of choices for isolation and scale to ATLAS 8 TeV data [212], with the cuts of eq. (5.0.1). As elsewhere, for both smooth-cone and (matched) hybrid isolation we use a cone of radius 0.4 and a threshold $E_T^{\text{thr.}} = 11 \text{ GeV}$, whilst for matched-hybrid isolation we use inner-cone

radius $R_d = 0.1$.

We begin in fig. 5.14 with the two fully-NNLO distributions $d\sigma/dM_{\gamma\gamma}$ and $d\sigma/d|\cos\theta_\eta^*|$. The features highlighted above can now be seen to dramatically improve the overall agreement of the prediction with the data.

We consider first the $M_{\gamma\gamma}$ distribution. The first panel shows that the overall prediction for the conventional scale choice and isolation procedure, $\mu_0 = M_{\gamma\gamma}$ with smooth-cone isolation, consistently underestimates the data by about 20%, except in the largest $M_{\gamma\gamma}$ bins. Agreement within the scale uncertainty band of the NNLO prediction occurs only at the extremes of the distribution, in the lowest and highest $M_{\gamma\gamma}$ bins.

The second panel shows that, in the low- $M_{\gamma\gamma}$ region, the agreement observed in the first panel is a direct consequence of the low- $M_{\gamma\gamma}$ enhancement for $\mu_0 = M_{\gamma\gamma}$ outlined previously. Without it, the suppression resulting from smooth-cone isolation prevents agreement in this region. Conversely, the third panel shows that without the additional suppressive behaviour of smooth-cone isolation on the low- $M_{\gamma\gamma}$ prediction, it grows substantially relative to the data, which does not follow the same low- $M_{\gamma\gamma}$ behaviour.

Comparing the first and third panels, we see that with $\mu_0 = M_{\gamma\gamma}$, moving from smooth-cone to hybrid isolation leads to a prediction in better agreement with the data, though still not consistently within the scale uncertainties of the theory calculation. We also see that with the scale choice $M_{\gamma\gamma}$, and without the suppression due to smooth-cone isolation, the low- $M_{\gamma\gamma}$ behaviour arising from the scale choice is untamed, and leads to a growing deviation between theory and data as $M_{\gamma\gamma}$ decreases.

The last panel shows that without either the enhancement due to $\mu_0 = M_{\gamma\gamma}$ for small $M_{\gamma\gamma}$, or the suppression in the same region due to smooth-cone isolation for small $\Delta R_{\gamma\gamma}$, we see agreement in this region between the theory prediction and the data. The combined effects on the overall normalisation of more permissive isolation and of the alternative scale choice $\mu_0 = \langle p_T^\gamma \rangle$ correct the 20% suppression throughout

the distribution, resulting in theory predictions and experimental measurements largely agreeing within the scale uncertainty bands throughout the distribution, except in the highest $M_{\gamma\gamma}$ bin where we might expect missing electroweak contributions to become significant.

We now turn to the $|\cos \theta_\eta^*|$ distribution, defined by

$$|\cos \theta_\eta^*| = \tanh \left(\frac{1}{2} |\Delta \eta_{\gamma\gamma}| \right) \quad (5.3.1)$$

which is plotted for reference in fig. 5.15. In the first panel in fig. 5.14 we see that the prediction with the scale choice $\mu_0 = M_{\gamma\gamma}$ and smooth-cone isolation substantially undershoots the data, by 15% at small rapidity-separations and 40% at high rapidity-separations. This is absent for the scale choice $\mu_0 = \langle p_T^\gamma \rangle$ in panels 2 and 4, and is therefore an artefact arising directly from the scale $M_{\gamma\gamma}$ and its approximately-exponential growth with rapidity separation as discussed in section 5.2.3. Any other scale that is independent of $\Delta y_{\gamma\gamma}$ (or, in the notation of section 5.2.4, with $\alpha = 0$) would be expected to show a similarly flat ratio to the data. Clearly, for fixed-order predictions made with $\mu_0 = M_{\gamma\gamma}$ to exhibit such a ratio, the PDFs would need to grow to counterbalance the suppression of the cross-section. It is not clear that this would be possible in such a way as to allow simultaneous agreement with data with both categories of scales.

As expected, between panels 1 and 3, and 2 and 4, the change in isolation between smooth-cone and hybrid-isolation yields an flat upwards normalisation, resulting in very good agreement across the rapidity range for the combination $\mu_0 = \langle p_T^\gamma \rangle$ and hybrid isolation.

5.3.2. Comparison to ATLAS data: two-way comparison

We have up to now separately investigated the effect of altering scale and isolation independently. Here we examine the combined effect on the agreement with ATLAS data of the simultaneous transition between the combinations corresponding to panels 1 and 4 of the plots in fig. 5.14, namely




Figure 5.14: Combined ratio plots for the four-way scales and isolation comparison, at NNLO (ratio to data).




Figure 5.15: Relationship between $|\cos \theta_\eta^*|$ and $\Delta\eta_{\gamma\gamma}$. The dashed grey lines indicate the relationship between ATLAS bins for $|\cos \theta_\eta^*|$ and the corresponding intervals for $\Delta\eta_{\gamma\gamma}$. At ATLAS, the experimental cuts $|y^\gamma| < 2.37$ restrict the rapidity separation to $|\Delta\eta_{\gamma\gamma}| < 4.74$. This cut only affects the result in the final bin, which otherwise extends to infinite rapidity separations.

- (a) $\mu_0 = M_{\gamma\gamma}$ with smooth-cone isolation, and
- (b) $\mu_0 = \langle p_T \rangle$ with hybrid isolation.

These are plotted for the six observables which ATLAS measured in fig. 5.16, with axis limits and layout set to enable easy comparison with the corresponding figure (fig. 5) in the ATLAS experimental paper [212].

Across all six distributions, combination (b) gives better agreement with data almost everywhere. The regions where agreement is notably worse are those in the neighbourhood of the Sudakov singularities described in section 4.2.4, and hence where poor agreement is expected in the absence of resummation. In these effectively-NLO distributions we continue to see an incomplete description of the data. We can infer from the SHERPA results of [212] that the missing radiative corrections that would feature in an NNLO diphoton-plus-jets calculation are required to adequately describe the data in these distributions.

For completeness, in fig. 5.17 we show the order-by-order breakdown of the




Figure 5.16: Analogue of Figure 5 from [212] showing the effects of the modified scale choice and isolation criteria on the prediction.

NNLO calculation for choice (b) of scale and isolation criterion, showing the relative magnitude of the NNLO corrections with these parameters.




Figure 5.17: Illustration of the perturbative convergence of the fixed-order predictions, for the six measured ATLAS distributions.

Electroweak corrections

The strong coupling constant is $\alpha_s(M_Z) \approx 0.118$, whilst the electromagnetic coupling constant is $\alpha_{\text{em}}(0) \approx 7.30 \times 10^{-3}$. As a consequence,

$$\alpha_{\text{em}}(0) \approx \alpha_s(M_Z)^{2.3}, \quad (6.0.1)$$

and so we might naïvely expect the first $\mathcal{O}(\alpha_{\text{em}})$ electroweak corrections to be of comparable magnitude to NNLO QCD corrections. Such corrections arise from the real radiation of final-state photons,¹ and virtual loop diagrams featuring the emission and reabsorption of a virtual boson (either a photon, or a massive gauge boson). In practice, as we have seen in chapters 4 and 5, the QCD corrections give K -factors larger than those that would naïvely be expected from perturbation theory, as new partonic channels become possible at higher orders, enhanced by the large small- x gluon PDF. Since the photon PDF is smaller by several orders of magnitude, the new photonic channels arising at NLO EW are PDF-suppressed rather than enhanced, so no large K -factors arise and the NLO EW corrections can be expected

¹ Since photons are massless, photonic real emissions must be included in order to cancel against the ε -poles of the corresponding $\mathcal{O}(\alpha_{\text{em}})$ virtual matrix elements, exactly as for QCD corrections in section 2.3.2.2. In contrast, loop matrix elements with virtual massive electroweak gauge bosons are separately finite, and the corresponding real radiation can be systematically chosen to constitute a separate process (e.g. $X + Z$ is entirely its own process, and not a radiative correction to the X final-state). This simplifies the book-keeping by preventing double-counting, (e.g. by avoiding the ambiguous distribution of these contributions among the real radiation of X for the Z final-state, Z for the X final-state, and the $X + Z$ final-state). Following [220, 221], we adopt this convention.

to be much less significant than the NNLO QCD corrections.

They are nevertheless important ingredients of a precision calculation, especially in the high-energy region in which large Sudakov double logarithms such as $\log^2(\hat{s}/M_Z^2)$ (arising from virtual soft and collinear gauge-bosons coupling to external legs) become dominant, when the scale of momentum transfers becomes much larger than the mass scales of the virtual particles in the loops. Because they scale with momentum transfer, these corrections are especially important as the centre of mass-energy of the collision increases and when collider luminosity allows distributional tails to be probed, including at the 13 TeV LHC. The leading global behaviour of these logarithms is universal, precisely as the corresponding ε -poles of section 2.3.2.1 are in QCD, and was calculated in [222, 223] and has been implemented in [224].

An alternative approach to applying the corrections approximately is to compute the electroweak corrections exactly, just as for the QCD corrections, by calculating the relevant matrix elements, with infrared divergences regulated as for QCD, and integrating them over the final-state phase space with specified fiducial cuts.

The electroweak corrections for diphoton production in association with up to two jets have previously been calculated in [225] using SHERPA and the one-loop matrix-element provider GoSAM [226, 227]. Here we present the calculation of the electroweak corrections to the diphoton process using antenna subtraction, in the approximation of neglecting the photon-initiated contributions. These have since been separately implemented [228] in the NNLOJET framework, allowing us to determine both the accuracy of the approximation, and the impact of the electroweak corrections on diphoton phenomenology.

6.1. Matrix elements and subtraction terms

The inclusion of electroweak corrections requires the modification of eq. (2.2.1) from an expansion in α_s to a double expansion in $(\alpha_{\text{em}}, \alpha_s)$,²

$$\begin{aligned} d\hat{\sigma}_{ab \rightarrow X} &= \left(\frac{\alpha_s}{2\pi}\right)^m \left(\frac{\alpha_{\text{em}}}{2\pi}\right)^n \sum_{i=0}^{\infty} \sum_{j=0}^{\infty} \left(\frac{\alpha_s}{2\pi}\right)^i \left(\frac{\alpha_{\text{em}}(\mu)}{2\pi}\right)^j d\hat{\sigma}^{(i,j)} \\ &= \left(\frac{\alpha_s}{2\pi}\right)^m \left(\frac{\alpha_{\text{em}}}{2\pi}\right)^n \left[d\hat{\sigma}^{\text{LO}} + \left(\frac{\alpha_s}{2\pi}\right) d\hat{\sigma}^{\text{NLO}} + \left(\frac{\alpha_s}{2\pi}\right)^2 d\hat{\sigma}^{\text{NNLO}} \right. \\ &\quad \left. + \left(\frac{\alpha_{\text{em}}}{2\pi}\right) d\hat{\sigma}^{\text{NLO EW}} + \mathcal{O}(\alpha_s^3, \alpha_{\text{em}}\alpha_s, \alpha_{\text{em}}^2) \right], \end{aligned} \quad (6.1.1)$$

where we can identify the QCD expansion at the leading order in α_{em} in the first line, and the lowest-order term of the electroweak expansion in the second.

In principle, it also requires the further modification [229] of the DGLAP equations of eq. (2.1.6), with the introduction of a photon PDF, with QED evolution equation

$$\begin{aligned} \frac{\partial f_\gamma(x, \mu_F^2)}{\partial \ln \mu_F^2} &= \frac{\alpha_{\text{em}}}{2\pi} \int_x^1 \frac{dz}{z} \left[\sum_i Q_{q_i}^2 \left[f_{q_i}\left(\frac{x}{z}, \mu_F^2\right) + f_{\bar{q}_i}\left(\frac{x}{z}, \mu_F^2\right) \right] P_\gamma^q(z) \right. \\ &\quad \left. + f_\gamma\left(\frac{x}{z}, \mu_F^2\right) P_\gamma^\gamma(z) \right], \end{aligned} \quad (6.1.2)$$

where

$$P_\gamma^\gamma(z) = -\frac{2}{3} \left[N_c \sum_i Q_{q_i}^2 + \sum_j Q_{\ell_j}^2 \right] \delta(1-z) + \mathcal{O}(\alpha_{\text{em}}), \quad (6.1.3)$$

and P_γ^q is as in eq. (4.1.4). Beyond $\mathcal{O}(\alpha_{\text{em}})$, further $\mathcal{O}(\alpha_s\alpha_{\text{em}})$ DGLAP terms arise in the evolution of QCD PDFs to account for QED-QCD mixing.

In this section we will neglect the photon-initiated contribution, formally by making the approximation that the photon PDF vanishes,

$$f_\gamma^p \equiv 0,$$

² For processes with, for example, four external quarks, the ‘Born’ process is not unique, which can lead to subtleties we do not dwell on here.




Figure 6.1: NNPDF3.1luxQED NNLO PDFs from [234], at 3 and 100 GeV, including the determination of the photon PDF f_γ^p . The noisy appearance relative to fig. 2.3 is ultimately due to the different ansätze used by the two collaborations to parametrise the PDFs.

and avoid calculating the photon-initiated matrix elements. In section 6.2.2 we will investigate the impact of this approximation. This assumption is *a priori* reasonable because fits of the photon PDF determined from experimental data [230–237], one of which is plotted in fig. 6.1, find it to be several orders of magnitude smaller than the other, QCD, PDFs, and hence well within the error of the calculation.

6.1.1. Real QED radiation

The triphoton matrix element for real photonic radiation from a quark line, $B_{2;3\gamma}^0$, is identical to $B_{3;2\gamma}^0$ of eq. (3.5.7), to give, with the above conventions,

$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{R EW}} = 8\pi^2 N_{q\bar{q}}^{\text{LO}} \frac{1}{3} Q_f^6 \delta_{f_2}^{f_1} \left[d\Phi_3 B_{2;3\gamma}^0 \right]. \quad (6.1.4)$$

An example Feynman diagram showing the $q\gamma$ -initiated crossing of $B_{2;3\gamma}^0$ is shown in fig. 6.2a. Since the matrix element is identical to the QCD matrix element, an A_3^0 -type antenna function will correctly subtract the soft- and collinear-divergences

precisely as in eq. (3.5.12), except that any of the photons can be unresolved,

$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{SII NLO EW}} &= 8\pi^2 N_{q\bar{q}}^{\text{LO}} \frac{1}{3} Q_f^6 \delta_{f_2}^{f_1} \\ &\times \left[d\Phi_3 \sum_{i=3}^5 A_3^0(\hat{1}_q, i_\gamma, \hat{2}_{\bar{q}}) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; j'_\gamma, k'_\gamma) \right], \end{aligned} \quad (6.1.5)$$

where j_γ and k_γ are in each case the indices of the two photons that are not i_γ .

In principle at this order there should be real-emission corrections in the $q\gamma$ -channel, analogous to eq. (3.5.16). However, since we are working in the approximation in which the photon-PDF is zero, these are not included.

6.1.2. Virtual electroweak corrections

We use the one-loop provider OPENLOOPS to provide the virtual matrix elements, which include divergent QED loop diagrams, such as the one shown in fig. 6.2b, as well as finite loop contributions with virtual massive vector bosons such as those shown in figs. 6.2c and 6.2d. The pole structure of the QED loop-matrix-elements is universal, as in eq. (2.3.28), and can be verified numerically to have the expected coefficients.

Integrated, the subtraction terms of eq. (6.1.5) give


$$\begin{aligned} d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{TII NLO}} &= N_{q\bar{q}}^{\text{LO}} Q_f^6 C_\varepsilon \delta_{f_2}^{f_1} \\ &\times \left[-d\Phi_2 \int \frac{dx_1}{x_1} \frac{dx_2}{x_2} \mathcal{A}_{q\bar{q},g}^0(x_1, x_2) B_{2;2\gamma}^0(\hat{1}'_q, \hat{2}'_{\bar{q}}; 3'_\gamma, 4'_\gamma) \right] \end{aligned} \quad (6.1.6)$$

and the QED mass-factorisation counterterm is


$$d\hat{\sigma}_{q_{f_1}\bar{q}_{f_2}}^{\text{MF NLO EW}} = N_{q\bar{q}}^{\text{LO}} Q_f^6 C_\varepsilon \delta_{f_2}^{f_1} \left[-d\Phi_2 \int \frac{dx}{x} \frac{2}{\varepsilon} p_q^{q(0)}(x) B_{2;2\gamma}^0 \right],$$

with p_q^q the colour-factor-stripped equivalent of the QCD splitting function eq. (2.1.7),


$$p_q^{q(0)}(z) = \left[\frac{1+z^2}{(1-z)_+} + \frac{3}{2} \delta(1-z) \right]. \quad (6.1.7)$$




(a) Photon-initiated crossing of QED real-radiation correction $B_{2;3\gamma}^0$.



(b) Example virtual diagram. Loop diagrams contribute by interfering with the tree-level diagram in which the quarks are necessarily the same flavour, so the virtual gauge boson can only be γ or Z .



(c) Example virtual diagram, which generates a pseudo-resonance at $M_{\gamma\gamma} \approx 2M_W$.



(d) Further examples of virtual electroweak correction diagrams.

Figure 6.2: Example Feynman diagrams contributing to the NLO electroweak corrections.

6.2. Results

Here we present the NLO electroweak results. We use a setup corresponding to the forthcoming release of ATLAS 13 TeV data [3],

$$p_T^{\gamma_1} > 40 \text{ GeV} \qquad p_T^{\gamma_2} > 30 \text{ GeV} \qquad (6.2.1a)$$

$$\Delta R_{\gamma\gamma} > 0.4 \qquad |y^\gamma| \in [0, 1.37) \cup (1.52, 2.37) \qquad (6.2.1b)$$

$$E_T^{\text{iso,part}} < 0.09 E_T^\gamma \qquad \text{within cone } \Delta R \leq 0.2. \qquad (6.2.1c)$$

We use the complex mass scheme³ [238, 239], and set $\alpha_{\text{em}}(0) = 1/137$ throughout as for the QCD calculation. For the comparison with the photon-initiated contributions, we use the NNPDF3.1luxQED PDF set [234], which contains non-zero photon PDFs.

For photon isolation in the presence of an additional photonic real emission, we generalise the above cuts to identify the third photon as a ‘long-distance’, resolved, photon if $p_T^{\gamma_3} \geq 30 \text{ GeV}$, and apply the $\Delta R_{\gamma\gamma} \geq 0.4$ cut to all identified photons.

6.2.1. Impact of electroweak corrections

The resulting EW corrections are shown in fig. 6.3 for a selection of differential cross-sections susceptible to the Sudakov behaviour at high momentum-transfers. They show a threshold effect at the Born threshold, a peak at $M_{\gamma\gamma} \approx 2M_W$ from the W -box diagrams, and the expected Sudakov logarithms in the tails. These pseudo-resonances lead to an overall-positive NLO correction, despite the negative corrections in the tails.

Their impact is consistently small, ranging from +10% at the Born $M_{\gamma\gamma}$ -threshold, to -10% in the high-energy limit, relative to the leading-order contribution. Overall,

³ In the complex-mass scheme, particle masses M_i are consistently replaced by the complex values μ_i , according to the formula

$$\mu_i^2 = M_i^2 - i\Gamma_i M_i. \qquad (6.2.2)$$

This regularises the divergent propagator $1/(p^2 - \mu_i^2)$ arising from on-shell internal massive gauge bosons, in a gauge-invariant way.




Figure 6.3: The effect of NLO electroweak corrections on the differential cross-sections, relative to leading-order, with respect to $M_{\gamma\gamma}$, p_T^1 and p_T^2 . The NNLO corrections are plotted for reference. The peaks correspond to $M_{\gamma\gamma} = 2p_T^1$ and $M_{\gamma\gamma} = 2M_W$.

they give a K -factor with respect to LO of approximately 1.02. Since the NLO and NNLO QCD K -factors are large, the impact of the corrections is extremely modest.

However, since they are in practice the NLO corrections only to the $q\bar{q}$ -channel, it remains likely that in the high-energy limit the $\mathcal{O}(\alpha\alpha_s)$ contributions (e.g. the leading-order electroweak corrections to the qg -channel) will be substantially larger, despite being formally suppressed by a further factor of α_s .

6.2.2. Impact of photon-initiated contributions

Finally, we examine the impact on the electroweak corrections of the approximation in which the photon PDF, already two orders of magnitude smaller than the quark- and gluon-PDFs, is set to zero. The different results are plotted in fig. 6.4. As expected, they are highly suppressed relative to the quark-antiquark contribution, and are only non-negligible in the regions of phase-space that are vetoed in the Born kinematics. This shows that the approximation of setting the photon-PDF to zero, and neglecting the photon-initiated contribution, is an insignificant one for phenomenology.




Figure 6.4: The NLO electroweak corrections, with and without the photon-initiated contributions, which can be verified to be negligible even with respect to the small $q\bar{q}$ -channel electroweak corrections. The non-analytic behaviour of the ratio is an artefact of the denominator crossing from being positive to negative, and is not of physical origin.

CHAPTER 7

LHC Run 2: outlook at 13 TeV

The preceding phenomenological sections of this thesis have focused on the application of the NNLOJET NNLO calculation to post-dicting observations made at ATLAS at 8 TeV, in Run I of the LHC. This data, based on an integrated luminosity of 20.2 fb^{-1} and released in 2016, remains the most recent and highest-quality data for diphoton production officially released by either experimental collaboration.

However, preliminary results representing 139 fb^{-1} of data collected by the ATLAS detector at 13 TeV in Run II of the LHC were presented at ICHEP 2020 and released preliminarily in [3]. We therefore apply the conclusions of the previous sections to this data.

As in the previous chapter, we use the ATLAS 13 TeV cuts from [3],

$$p_T^{\gamma_1} > 40 \text{ GeV} \qquad p_T^{\gamma_2} > 30 \text{ GeV} \qquad (7.0.1a)$$

$$\Delta R_{\gamma\gamma} > 0.4 \qquad |y^\gamma| \in [0, 1.37) \cup (1.52, 2.37) \qquad (7.0.1b)$$

$$E_T^{\text{iso,part}} < 0.09 E_T^\gamma \qquad \text{within cone } \Delta R \leq 0.2. \qquad (7.0.1c)$$

The important changes with respect to the 8 TeV ATLAS cuts of eq. (5.0.1) are to the isolation criteria, with a smaller cone radius of 0.2 (vs 0.4), and with a variable E_T^{iso} , with $\varepsilon_\gamma = 0.09$ and $E_T^{\text{thr.}} = 0$ (vs. $\varepsilon_\gamma = 0$, $E_T^{\text{thr.}} = 11 \text{ GeV}$). Based on the conclusions of chapters 4 and 5 we can therefore expect that the impact of moving

between smooth-cone and hybrid isolation is likely to be less substantial for the 13 TeV cuts than it was at 8 TeV, since the region of phase-space on which radiation is unphysically suppressed by smooth-cone isolation is substantially smaller.

The NNLO predictions using $\mu_0 = \langle p_T^\gamma \rangle$ and hybrid isolation are shown in figs. 7.1 and 7.3 alongside the alternative scale and isolation prescriptions, and the preliminary ATLAS data. Figure 7.1 shows the data for the two observables closest¹ to those considered at 8 TeV in chapter 5, whilst fig. 7.3 shows two fully-NNLO observables measured for the first time, the transverse momenta of the two photons.

All four distributions show excellent agreement with the data for this choice, whilst the scale $M_{\gamma\gamma}$ again shows the expected deviations at low- $M_{\gamma\gamma}$ and high $|\cos \theta_{\text{CS}}^*|$. Away from the low-energy and high-energy extremes of the $d\sigma/dM_{\gamma\gamma}$ distribution, the measured data lies outwith the scale uncertainty bands only at $M_{\gamma\gamma} \approx 80 \text{ GeV}$, the Born threshold, because of non-analyticity at the Sudakov critical point. This is shown in detail in fig. 7.2.

As noted previously, the scale uncertainty bands of the $M_{\gamma\gamma}$ distribution are in general much smaller for the scale choice correlated with the distribution than for the uncorrelated choice $\langle p_T^\gamma \rangle$, except in the pathological limit $M_{\gamma\gamma} \rightarrow 0$. This leads to an apparent tension between theory and experiment for $\mu_0 = M_{\gamma\gamma}$ for the majority of the distribution ($M_{\gamma\gamma} \gtrsim 200$) that is only resolved by changing the functional form of the scale. This can also be seen in the high-rapidity-separation limit of the $|\cos \theta_{\text{CS}}^*|$ distribution.

As expected, the suppressive effect of smooth-cone isolation is substantially less severe for $R = 0.2$ than it was for $R = 0.4$, and so the effect of using hybrid isolation

¹ The scattering angle [240] with respect to the beam axis in the Collins-Soper frame [241] $|\cos \theta_{\text{CS}}^*|$ was introduced to improve the experimental resolution of collider measurements by relying primarily on angular measurements, which can be made precisely, rather than energy measurements, which are subject to greater uncertainties. It can be defined in a frame-invariant way through

$$|\cos \theta_{\text{CS}}^*| := \frac{2p_T^{\gamma_1} p_T^{\gamma_2} \sinh \Delta y_{\gamma\gamma}}{M_{\gamma\gamma} \sqrt{M_{\gamma\gamma}^2 + (p_T^{\gamma\gamma})^2}} \equiv \frac{\sinh \Delta y_{\gamma\gamma}}{\sqrt{\cosh \Delta y_{\gamma\gamma} - \cos \Delta \phi_{\gamma\gamma}}} \left[\cosh \Delta y_{\gamma\gamma} + \frac{1}{2} \left(\frac{p_T^{\gamma_1}}{p_T^{\gamma_2}} + \frac{p_T^{\gamma_2}}{p_T^{\gamma_1}} \right) \right]^{-\frac{1}{2}}, \quad (7.0.2)$$

so, as for $|\cos \theta_\eta^*|$, the high-rapidity-separation region is mapped onto $|\cos \theta_{\text{CS}}^*| \approx 1$.




Figure 7.1: NNLO predictions for ATLAS at 13 TeV, for observables most closely corresponding to those studied in detail at 8 TeV. ATLAS data extracted manually from the plots of [3] (experimental errors not included).




Figure 7.2: The LO, NLO and NNLO predictions for the $d\sigma/dM_{\gamma\gamma}$ distributions in the neighbourhood of the Born kinematics threshold, $M_{\gamma\gamma} \approx 80$ GeV. The visual smoothness of the NNLO distribution is limited by the small bin-size and the Monte Carlo statistics, but the important features remain visible. The distortion induced by the breakdown of perturbation theory at the Sudakov singularity at the boundary of the Born phase space causes the data in this bin to lie outwith the theory scale-uncertainty.

instead is smaller than in chapter 5. Nevertheless some of the features remain apparent, including the suppressive effect of smooth-cone isolation at low- $\Delta R_{\gamma\gamma}$, visible in the $M_{\gamma\gamma}$ distribution but no longer large enough to fully compensate for the growth of $\alpha_s(M_{\gamma\gamma})$ in this limit.

The two photon- p_T distributions shown in fig. 7.3 show excellent agreement between theory and data. As for the other distributions, the twofold scale variation about $\mu_0 = M_{\gamma\gamma}$ fails to encompass the data, but because here the observable and the scale are uncorrelated the deviation is reduced to a small global normalisation.

Finally, we present combined NNLO QCD + NLO EW predictions for the preliminary 13 TeV ATLAS data, in fig. 7.4. This is the first calculation at this order. We see that the effect of the electroweak corrections is modest, as expected, and makes negligible difference to the description of the data.

We can therefore conclude that the prospects for precision diphoton phenomenology at higher centre-of-mass energies are excellent. Whilst the scale uncertainties as estimated by the conventional scale-variation procedure remain substantial, due to its large sensitivity to the leading-order gg -channel contribution, the calculation gives excellent agreement with data. The recent calculation of the five-parton, two-loop matrix-elements should soon allow this level of precision to be extended beyond the Born kinematics to the distributions which are trivial at leading-order, including the transverse momentum of the diphoton system $p_T^{\gamma\gamma}$ and the azimuthal angular separation $\Delta\phi_{\gamma\gamma}$.




Figure 7.3: NNLO predictions for ATLAS at 13 TeV. ATLAS data extracted manually from the plots of [3] (experimental errors therefore not included).




Figure 7.4: LO, NLO and NNLO QCD, and NNLO QCD + NLO EW, predictions for ATLAS at 13 TeV. ATLAS data extracted manually from the plots of [3] (experimental errors therefore not included).




Figure 7.4: LO, NLO and NNLO QCD, and NNLO QCD + NLO EW, predictions for ATLAS at 13 TeV. ATLAS data extracted manually from the plots of [3] (experimental errors therefore not included).




Figure 7.4: LO, NLO and NNLO QCD, and NNLO QCD + NLO EW, predictions for ATLAS at 13 TeV. ATLAS data extracted manually from the plots of [3] (experimental errors therefore not included).

Conclusions

In this thesis, we have outlined the application of antenna subtraction to the calculation of direct diphoton production, both at next-to-next-to-leading order in QCD, and at next-to-leading-order in the electroweak theory. This is the first NNLO calculation with a fully-local subtraction scheme, and therefore unaffected by the possible influence of photon isolation on the accuracy of other, non-local, subtraction schemes.

We have implemented this in the Monte Carlo parton generator NNLOJET, and used it to study the theoretical and phenomenological consequences of two unphysical choices that must be made, in order to generate predictions for comparison to data: the choice of theoretical photon isolation procedure, which can at best approximate the criterion applied experimentally, and the choice of renormalisation and factorisation scale.

In both cases, we found that the uncertainties implicit in making the conventional choices (of smooth-cone isolation and $\mu_0 = M_{\gamma\gamma}$) are substantial enough to be dominant sources of theory uncertainty, and are underestimated by the conventional parameter variation of scale coefficients and isolation parameters.

We have studied reasonable alternatives to these choices, which have not previously been applied to diphoton production, and made a detailed study of an altern-

ative combination, choosing the average transverse momentum of the two photons for μ_0 , and hybrid isolation. The systematic investigation of the consequences of moving between the four possible combinations of choices revealed that an interplay between the effects of the conventional choices on the same regions of phase-space disguises their impact. We have applied these conclusions to ATLAS 8 TeV data, and concluded that, far from failing to describe the data, it is possible for an NNLO QCD calculation to describe the data very well.

We have then turned to the application of antenna subtraction to QED subtraction, to allow the calculation of electroweak corrections. This is especially simple for the diphoton process, and was implemented in the NNLOJET framework and used to study the relative significance of these corrections for phenomenology at 13 TeV.

Finally, the NNLO QCD and NLO electroweak calculations were combined and applied to the prediction of newly-released 13 TeV ATLAS data from Run II of the LHC, the first calculation at this order. The conclusions formed from the 8 TeV study continue to apply at 13 TeV, with excellent agreement between theory and data for hybrid isolation and $\mu_0 = \langle p_T^\gamma \rangle$.

Set in context, this work represents the first successful description of ATLAS data for the diphoton process with precision theory calculations. This highlights the importance of fully understanding theory uncertainties when making precision calculations, and serves to caution against attempts to estimate variation over a function space purely through a one- or two-parameter subspace. This will be an important lesson as LHC Run III begins, narrowing still further the target precision for theory calculations.

APPENDIX A

Validation and testing in NNLOJET

New processes implemented in NNLOJET undergo a stringent series of tests to verify that the matrix elements, typically complex assemblies of individual spinor-helicity amplitudes, have been correctly coded from mathematical formulae, and that the antenna subtraction terms implemented do indeed subtract their divergent limits.

This new implementation of the diphoton process has been subjected to these tests, which we summarise here. The ordering of these tests is systematic: in general, the diagnostic power of later tests depends on the prior verification of elements tested independently in earlier tests.

A.1. Pointwise validation of matrix elements

All tree-level matrix elements, the one-loop, five-parton matrix elements $B_{3;2\gamma}^1$, $\tilde{B}_{3;2\gamma}^1$, and the one-loop gluon-gluon box matrix element have been implemented from scratch in native **Fortran** code.

These have been tested, in all crossings, against fully-automated matrix elements generated by integrand reduction within MADGRAPH [133] by MADLOOP [242], for a number of phase-space points randomly generated by RAMBO [243]. Sample output is shown in listings A.1 and A.2. Loop matrix elements are tested at each relevant level of the ε -expansion, as shown in listing A.2.

```

Matrix element = 2.0803456038366773E-011 GeV^-4

LCol 2.0683216882033166E-009
SLCol -7.1189908520103213E-011
NNLOjet matrix element = 2.0803456038366808E-011 GeV^-4
Ratio: NNLOjet over MadGraph 1.0000000000000018
-----
Percentage error: 1.7085102052316636E-013 %
-----

Average ratio: 1.0000000000000024
Min ratio: 0.99999999999999423
Max ratio: 1.00000000000000122
Average relative error: 2.4424906541753444E-015

```

Listing A.1: The tail of point-test output comparing hard-coded NNLOJET matrix-elements with those generated automatically by MADGRAPH, here testing $B_{4;2\gamma}^0$ and $\tilde{B}_{4;2\gamma}^0$. The average, minimal and maximal ratios are computed over a customisable number of randomly-generated phase-space points, here 20.

Beyond this, the OPENLOOPS interface used for the electroweak corrections in chapter 6 has also been used to validate the matrix elements for the photonic real-emission directly within NNLOJET, by performing identical phase-space integrations with identical VEGAS initial seeds for a specific partonic channel, with only the origin of the matrix element (and the associated auto-generated factors) switched between the two runs. Because the VEGAS warmup is iterative, it amplifies even small differences between matrix elements. Getting the same result within machine precision after many warmup iterations is therefore a powerful check on the agreement between the two.

```

BORN 1.3927726547794248E-006
SINGLE POLE (MadFKS) -2.8357518937063263E-007
DOUBLE POLE (MadFKS) -1.4822110171280097E-007

SINGLE POLE (MadLoop) -2.8357518937062617E-007
DOUBLE POLE (MadLoop) -1.4822110171280039E-007
FINITE PART (MadLoop) 1.4661106444095933E-007

BORN (NNLOJET) 1.3927726547794231E-006
SINGLE POLE (NNLOJET) -2.8357518937063226E-007
DOUBLE POLE (NNLOJET) -1.4822110171280073E-007
FINITE PART (NNLOJET) 1.4661106444094972E-007

NJ/MG: born 0.99999999999999878

NJ/MG: double 1.0000000000000022
NJ/MG: single 1.00000000000000215
NJ/MG: finite 0.99999999999993450

```

Listing A.2: Extract from point-testing output for virtual matrix elements $B_{3;2\gamma}^1$ combined with $\tilde{B}_{3;2\gamma}^1$ and $\overline{B}_{3;2\gamma}^1$. The contribution from the latter can be tested separately by varying n_f . The different levels of the ε -expansion are each checked independently.


A.2. Validation of subtraction terms

A.2.1. Cancellation of real IR divergences


It is crucial that the subtraction terms, as implemented, correctly mimic the divergent limits of the real-radiation phase-space. To test this we consider their ratio in each relevant limit, as a function of how far into the limit we probe. These ratios are binned as histograms and plotted for a number of randomly-generated phase-space points, as shown in fig. A.1. If the subtraction term is correct, the points will cluster into the neighbourhood of 1 as the invariants approach the exact limit.

A.2.2. Cancellation of virtual ε -poles


As outlined in detail in section 3.5, in a correct implementation of the virtual matrix elements and the antenna subtraction counterterms, the coefficients of the ε -poles necessarily cancel. Since the pole pieces of the matrix elements are separately



(a) Example spike test for $q\bar{q} \rightarrow gg\gamma\gamma$, testing the soft-gluon limit.



(b) Example spike test for $gg \rightarrow q\bar{q}\gamma\gamma$, testing the double-collinear limit in which the quark and antiquark each become collinear to an initial-state gluon.



(c) Example spike test for $q\bar{q} \rightarrow Q\bar{Q}\gamma\gamma$, testing the double-soft limit in which the final-state quark-antiquark pair simultaneously become soft

Figure A.1: Example spike-tests for double-real matrix elements.
The x -parameter determines the relative size of the relevant invariant $s_{i\dots k}$; the smaller it is, the further the test probes into the divergent limit.

validated against the automated implementation of MADLOOP, and the unintegrated subtraction terms are separately validated against the real-radiation matrix elements, this ensures that the separate components are compatible.

This cancellation is checked

- analytically for the two-loop matrix elements, using `Form` within a `Maple` routine on the `Maple` antenna-subtraction source files, and
- numerically for all loop-levels within the NNLOJET program at runtime (with the ‘pole-check’ mode selected in the NNLOJET runcard).

In the latter case, in the event of the test failing, it can be applied separately to individual partonic channels and colour levels to identify the source of the error.

A.2.3. Finite term consistency checks

As described in chapter 3, subtraction terms that correctly cancel the divergent limits of the matrix element are not unique: only their behaviour in each divergent limit is fixed. As a result, if one subtraction term passes the tests, any other term that differs by a regular function will likewise pass the tests, and all such subtraction terms will return the same ε -pole coefficients upon integration.

The spike- and pole-tests only test that the divergent parts of the subtraction terms and matrix elements cancel appropriately, and not that the finite remainders are consistent between integrated and unintegrated subtraction terms. An inconsistency would prevent the numerical integration over the unintegrated subtraction terms from cancelling as intended against the integrated-subtraction-terms.

This is mitigated against through ‘layer-checks’. These are `Maple` tests which compare the `Maple` subtraction terms contributing to $d\hat{\sigma}_{ab}^S$, $d\hat{\sigma}_{ab}^T$ and $d\hat{\sigma}_{ab}^U$ against each other (and $d\hat{\sigma}_{ab}^{S\text{ NLO}}$ against $d\hat{\sigma}_{ab}^{T\text{ NLO}}$ at NLO). They impose a set of symbolic identities, also coded in `Maple`, to verify that these terms do indeed cancel as intended.

A.3. Integrated cross-section tests

As is clear from section 3.5, the colour, charge, and symmetry factors associated with each matrix element are not completely straightforward to derive and verify, and it is easy to make mistakes computing them individually. For this reason, they are generated automatically for each NNLOJET process by `Maple` scripts accounting for colour and symmetry factors.

The phase-space generator and integration routines, and the PDF interface for convolution with parton distribution functions provided by LHAPDF [173], are common to several NNLOJET processes and have undergone extensive testing. We therefore assume their validity here.

A.3.1. Technical cut dependence

As discussed in section 2.6.2.2, in practice the divergent limits cannot be integrated over numerically, as the cancellation between divergent matrix element and divergent subtraction term, although theoretically exact, is subject to a loss of floating-point precision which can leave large remainders.

These regions are avoided through the use of a technical cut, which prevents any invariant from becoming too small by imposing

$$\min_{i,j} s_{ij} \geq t \hat{s}_{12}. \quad (\text{A.3.1})$$

Phase-space points failing this condition are discarded.

Whether or not the matrix elements contain a divergent limit as $s_{ij} \rightarrow 0$, we know from the spike-tests that the subtraction-terms render the integrand finite. The error induced by excluding this region of phase-space should therefore be proportional to t .

In the presence of imperfectly-cancelled divergences, however, t itself acts as the regulator of the divergent limit, leading to a strong dependence of the cross-section on t . Examining this dependence therefore allows us to detect, indirectly,

imperfectly-cancelled divergences.

Beyond this, t -dependence has practical consequences for precision calculations. From the above discussion it is clear that t cannot be set too small without risking numerical miscancellations between divergent terms, and cannot be set too large without risking excluding a non-trivial portion of phase-space. For practical purposes, therefore, we compute a value for t at which the t -dependence plateaus to within statistical uncertainties. For the results in this thesis we have used $t = 10^{-8}$.

A.3.2. Scale variation

The theoretical dependence of the QCD cross-section on the renormalisation scale, derived in section 2.4.1, can be used to test the calculation, by ensuring that the numerical results obtained by running NNLOJET for different scale choices reproduce the expected analytical scale-dependence of the cross-section.

This is shown in fig. A.2. If the test is initially failed, it can be applied to individual colour-levels and partonic channels to help diagnose the source of the error.

A.4. Validation against other codes

In addition to internal checks, which can be applied to NNLOJET without requiring external comparison, we have compared the results against those obtained by other programs.

A.4.1. LO and NLO

As for the general testing program, these tests can be ordered systematically to allow bugs to be isolated to a single part of the calculation, exploiting the fact that the ingredients of the NNLO diphoton calculation also contribute to LO and NLO calculations of diphoton production in association with up to two jets. This is shown in table A.2. The Born, real, and double-real matrix elements can be tested




Figure A.2: Scale-variation test for the diphoton calculation. The blue and red lines and data points are extrapolated from the $\mu_R = 90$ GeV result using the renormalisation group equations eq. (2.4.3) at NLO and NNLO respectively. The NNLOJET results can be seen to have the expected overall dependence upon μ_R , distributed across the matrix elements and partonic channels.

directly against a cross-section with equivalent setup calculated by an automated leading-order code (in this case, against both MATRIX and SHERPA).

Once the B, R and RR components have been validated, they can be combined with the V and RV components and assembled into LO calculations of the diphoton and the diphoton-plus-jet process, and again validated against automated implementations (again, against both MATRIX and SHERPA). The separate prior validation of the Born, real and double-real contributions ensures that discrepancies arising at this step are due either to the virtual matrix elements, or to the subtraction procedure.

Only at this point is the code ready for fully-NNLO validation, which is typically the most computationally expensive step.

	4-parton	5-parton	6-parton
tree	B	R	RR
one-loop	V	RV	
two-loop	VV		

Table A.1: Diagrammatic representation of the parts of an NNLO calculation, and the lower-order calculations through which the constituent parts can be tested. Here red represents the $\gamma\gamma$ calculation, purple the $\gamma\gamma + j$ calculation, and blue the $\gamma\gamma + 2j$ calculation. Dot-dashed lines represents LO of the corresponding calculation, and dashed lines NLO.

A.4.2. NNLO

The NNLOJET implementation of diphoton production is the third independent NNLO calculation, following 2γ NNLO [202], MCFM [203] and the incorporation of the two-loop matrix element code from 2γ NNLO into the more recent Monte Carlo framework MATRIX [210].

At NNLO, the 2γ NNLO and MATRIX calculations use q_T -subtraction, whilst MCFM uses N -jettiness slicing. This calculation is therefore the first to use a fully local subtraction scheme.

This leads to potential problems for precision validation, since as discussed in section 2.6.2.1, cross-sections computed using q_T - or N -jettiness subtraction only give the exact result in the limit of vanishing technical cut $\tau = t \rightarrow 0$. For non-zero cut-parameter ($r_{\text{cut}} = \min p_T^{\gamma\gamma}/M_{\gamma\gamma}$ for q_T -subtraction, or τ_{cut} for N -jettiness), the result is potentially subject to power corrections, which may be problematic for validation despite being too small to be relevant for phenomenology.

This is especially important for processes containing isolated photons, for which the q_T -subtraction r_{cut} -dependence has been found [210] to be uncommonly large, as shown in fig. A.3. This is also expected to be true for N -jettiness subtraction, due to the enhanced power corrections arising from fiducial cuts on the final-state




Figure A.3: Illustration of the polynomial extrapolation of $\sigma(r_{\text{cut}})$ to $r_{\text{cut}} = 0$, as performed by MATRIX, and the resulting uncertainty. The absence of a plateau indicates the residual cut-dependence. Taken from [210].

photons and photon isolation cuts [244].

To address the slower convergence of q_T -subtraction for processes with isolated photons, the authors of MATRIX recommend using $r_{\text{cut}} = 0.05\%$ for such processes rather than the default of $r_{\text{cut}} = 0.15\%$ where the target precision is 0.5% or better, to better approximate the $r_{\text{cut}} \rightarrow 0$ limit [210]. The additional uncertainty introduced in the (automated) extrapolation into this limit accounts for the increase in numerical error from ± 10 fb at $r_{\text{cut}} = 0.05\%$ to ± 160 fb in the $r_{\text{cut}} \rightarrow 0$ limit.

We follow this recommendation in the validation below, and further compare distributions, which are not extrapolated, in addition to cross-sections.

A.4.2.1. Fiducial cross-sections

Table A.2 summarises the cross-sections obtained at LO, NLO and NNLO for each of the available NNLO calculations, with an identical setup. The results for 2γ NNLO and MCFM have been taken from their respective papers [202, 203]; those for MATRIX and NNLOJET have been computed using identical PDFs and fiducial cuts, with $\sqrt{s}$ of 14 TeV and fiducial cuts of

$$p_T^{\gamma_1} > 40 \text{ GeV} \qquad p_T^{\gamma_2} > 25 \text{ GeV} \qquad (\text{A.4.1a})$$

$$M_{\gamma\gamma} \in [20, 250] \text{ GeV} \quad |y^\gamma| \in [0, 2.5) \quad (\text{A.4.1b})$$

$$\text{Smooth isol.: } E_{\text{T}}^{\text{iso}} = 0.5 p_{\text{T}}^\gamma \quad \text{within cone } \Delta R \leq 0.4, \quad (\text{A.4.1c})$$

with isolation profile function eq. (4.2.6) with $n = 1$,

$$\chi(r; R) = \left(\frac{1 - \cos r}{1 - \cos R} \right), \quad (\text{A.4.2})$$

and $\alpha = 1/137$.

As in the reference calculations we use the MSTW 2008 PDFs at the appropriate order in each case, and take for our central renormalisation and factorisation scales the invariant mass of the diphoton pair $\mu_{\text{F}} = \mu_{\text{R}} = M_{\gamma\gamma}$, with the scale-dependence of the calculation tested through the three-point scale variation $\mu_{\text{F}} = \mu_{\text{R}} = 2M_{\gamma\gamma}$ and $\mu_{\text{F}} = \mu_{\text{R}} = \frac{1}{2}M_{\gamma\gamma}$.

Code	Fiducial cross-section σ [fb]		
	LO	NLO	NNLO
2 γ NNLO [202]	$5712 \pm 2 \begin{smallmatrix} +607 \\ -667 \end{smallmatrix}$	$26402 \pm 25 \begin{smallmatrix} +357 \\ -179 \end{smallmatrix}$	$40269 \pm 250 \begin{smallmatrix} +1368 \\ -1969 \end{smallmatrix}$
MCFM [203]	$5710 \pm 1 \begin{smallmatrix} +605 \\ -667 \end{smallmatrix}$	$26444 \pm 12 \begin{smallmatrix} +334 \\ -134 \end{smallmatrix}$	$40453 \pm 30 \begin{smallmatrix} +1611 \\ -2232 \end{smallmatrix}$
MATRIX	$5714 \pm 1 \begin{smallmatrix} +607 \\ -668 \end{smallmatrix}$	$26475 \pm 3 \begin{smallmatrix} +339 \\ -130 \end{smallmatrix}$	$40477 \pm 10 \begin{smallmatrix} +1670 \\ -2148 \end{smallmatrix}$
			$40184 \pm 160 \begin{smallmatrix} +1635 \\ -2111 \end{smallmatrix}$
NNLOJET	$5712 \pm 1 \begin{smallmatrix} +607 \\ -668 \end{smallmatrix}$	$26474 \pm 7 \begin{smallmatrix} +339 \\ -130 \end{smallmatrix}$	$40328 \pm 22 \begin{smallmatrix} +1649 \\ -2125 \end{smallmatrix}$

Table A.2: Validation results for the central scale $\mu_{\text{F}} = \mu_{\text{R}} = M_{\gamma\gamma}$, accompanied by the associated statistical error for the central scale, and asymmetric scale uncertainties from the three-point scale variation. The first MATRIX NNLO result given is for finite $r_{\text{cut}} = 0.05\%$, and the second is the result of MATRIX's automatic quadratic extrapolation $r_{\text{cut}} \rightarrow 0$. The MCFM result is an extrapolation from runs at finite τ_{cut} , as described in [203].

All four calculations agree at LO, and the MATRIX and NNLOJET results agree at NLO. The 2γ NNLO and MCFM results deviate from both MATRIX and NNLOJET by approximately 0.3% and 0.1% respectively.

The diphoton process within MATRIX is fully automated at NLO, using the automated one-loop amplitude provider OPENLOOPS for its tree-level and one-loop amplitudes, whilst 2γ NNLO and MCFM use hand-written `C++` and `Fortran` routines implementing algebraic amplitudes as code.

At NNLO, the MATRIX and NNLOJET computations agree within the target precision of 0.5%, and all agree within 1%. All lie within the reported errors of the 2γ NNLO calculation in the literature, and all lie between the extrapolated and unextrapolated MATRIX cross-sections.

A.4.2.2. Differential distributions

Finally we compare the differential cross-sections produced by MATRIX with those produced by NNLOJET with the above cuts, for a range of differential cross-sections, in fig. A.4.

The agreement can be seen to be comfortably per-mille for LO and NLO, and substantially smaller than 1% at NNLO.

A.5. Regression tests

To prevent errors from being introduced into previously-validated calculations by subsequent changes to the NNLOJET code, a nightly system of bespoke regression tests is run to verify output from the latest development version of NNLOJET against that obtained with earlier versions.

These are based on `Python` routines, including one which automatically downloads and compiles the latest version of the NNLOJET source code committed to the private repository, another which acts as a wrapper for the execution of NNLOJET with each of a list of runcards, and an analysis routine which compares the output




Figure A.4: Ratio of differential distributions as calculated by NNLOJET and by MATRIX, with $\mu_R = \mu_F = M_{\gamma\gamma}$ and the setup as above. The statistical uncertainty is only substantial in the tails of distributions, where the cross-section is exponentially suppressed. Agreement can be seen to be comfortably within one part in 1000 at NLO, and within 1% at NNLO, though limited by statistical uncertainties in the tails of distributions.

to reference output. Where there are non-trivial changes to the output, the test is deemed to have failed, and the failing runcard is identified to the authors in an email sent upon completion.

A.6. Implementation and validation of soft-drop

The soft-drop isolation criterion introduced in section 4.2.3.1 has been implemented from scratch in NNLOJET, alongside soft-drop reclustering more generally.

A.6.1. Implementation

The implementation of soft-drop is designed to be minimally intrusive to the remaining NNLOJET code. This is slightly delicate, since in general the photon isolation and jet clustering functions are independent of one another (and contained in separate Fortran modules `PhotonIsolation_mod` and `EvalObs_mod` each with private routines), with each acting separately on the underlying partonic momenta and not requiring data from the other. For soft-drop isolation, which uses jet information for photon isolation, they must be intertwined.

Typically, isolation is applied before jet clustering, to save on the computational cost of jet clustering in events that will subsequently be vetoed by isolation cuts. In the soft-drop case, the jet clustering routine `cluster_jet` is instead first called as part of the isolation procedure, with the results, `ClusterHist_jet`, stored into a new variable `ClusterHist_orig` inside a dedicated (thread-private) Fortran module `Softdrop_mod`. For final-states in which there are n particles subject to the jet algorithm, `ClusterHist_jet` is an n -dimensional vector of integers such that

$$\text{ClusterHist_jet}(j) = i \tag{A.6.1}$$

if protojet¹ j was merged into protojet i .

¹ All particles subject to the jet algorithm are initially labelled ‘protojets’; as they are clustered, fewer and fewer protojets remain. After the clustering has been performed, the protojets are compared to the jet cuts, to determine which protojets constitute ‘jets’; their properties are then saved into

The photon isolation routine then calls the subroutine `recluster_jet`, which reconstructs the partonic constituents of each jet from the clustering history and applies the jet clustering algorithm again only to those particles, with the appropriate choice of generalised k_T -algorithm (typically Cambridge–Aachen for soft-drop isolation, but this is fully general and can be specified in the NNLOJET runcard).

At each protojet-merging in the reclustering step, the soft-drop condition of eq. (2.5.6),

$$\frac{p_T^2}{p_T^1 + p_T^2} > z_{\text{cut}} \left(\frac{\Delta R_{12}}{R} \right)^\beta, \quad (\text{A.6.2})$$

is tested, and the result (pass or fail) stored alongside the indices of the protojets in a $3 \times n$ array, `jpass`. This is generalisable to any condition that could be applied at the merging step. Each merging is stored in `ClusterHist_jet` as before, but with softer protojets always merged into harder protojets (i.e. $p_T^j \leq p_T^i$ above, rather than the default ‘nested-do-loop’ ordering $i < j$). The final clustering tree is then traversed backwards. The last row of `jpass` that has the jet index in the first column and a 1 in the last, indicating that the merging passed soft-drop condition, gives the final merge of the soft-dropped jet. The partons are then reclustered according to the `jpass` data up to that point, leaving jet arrays within NNLOJET exactly as they would have been if only one pass of the clustering algorithm had occurred (except that jets may contain photons), to allow other jet-defined functions and observables to work correctly.

The photon isolation routine then converts these results, which so far correspond simply to soft-drop reclustering, into soft-drop isolation, by testing each registered final-state photon i_γ to ensure that

- (i) the photon remains within its own protojet, i.e. `clusterHist_jet`(i_γ) = i_γ ,
- (ii) no other particles have been merged into its protojet, i.e.

$$\text{clusterHist_jet}(j) \neq i_\gamma \quad \forall j \neq i_\gamma. \quad (\text{A.6.3})$$

a dedicated array of `Jet`-type objects, `jets`, sorted in order of descending transverse momentum.

If either of these tests is failed for any final-state photon, it is not isolated and the event is rejected.² If they are passed, the event contains two isolated final-state photons, as determined by the soft-drop procedure. The photon-jets are then removed from the `jets` array, to allow the jet-observables to be used as normal.

A.6.2. Validation

The soft-drop reclustering routine described above has been validated against the `C++` code `FASTJET`, using a bridge code to apply `FASTJET`'s soft-drop routine directly within `NNLOJET` as an alternative backend. For this to be enabled, `NNLOJET` must be compiled against `FASTJET` (with the `RecursiveTools` library installed) using `jet=fastjet` as an argument to the `make` command.

The testing routines apply both the native `NNLOJET` and the `FASTJET` alternative algorithms to the same partonic momenta, and compare the returned groomed jets, printing a warning message if they differ. This has been run for several very large testing runs (with matrix elements turned off for computational efficiency), including with cuts chosen to drive the partonic momenta into regions of phase-space expected to be pathological, with identical results found in all cases.

² Here we consider exclusively QCD corrections, so every parton-level photon must be isolated for the event to pass the process-defining cuts. For real electroweak corrections, one parton-level photon could be clustered into a jet without incurring a veto.

Bibliography

- [1] T. Gehrmann, E.W.N. Glover, A. Huss and J. Whitehead, *Scale and isolation sensitivity of diphoton distributions at the LHC*, *JHEP* **01** (2021) 108 [2009.11310].
- [2] S. Amoroso et al., *Les Houches 2019: Physics at TeV Colliders: Standard Model Working Group Report*, in *11th Les Houches Workshop on Physics at TeV Colliders: PhysTeV Les Houches*, Mar., 2020 [2003.01700].
- [3] ATLAS collaboration, *Measurement of the production cross-section of isolated photon pairs in pp collisions at 13 TeV with the ATLAS detector*, Tech. Rep. ATLAS-CONF-2020-024, CERN, Geneva (July, 2020).
- [4] L. O’Raifeartaigh, *Group Structure of Gauge Theories*, Cambridge University Press (1988).
- [5] R. Gilmore, *Lie groups, physics, and geometry*, Cambridge University Press (2008).
- [6] E. Cartan, *Sur la structure des groupes de transformations finis et continus*, Ph.D. thesis, Université de Paris, 1894.
- [7] A. Sommerfeld, *Zur Quantentheorie der Spektrallinien*, *Annalen der Physik* **356** (1916) 125.
- [8] W. Greiner and J. Reinhardt, *Field quantization*, Springer (1996).

-
- [9] M.E. Peskin and D.V. Schroeder, *An Introduction to Quantum Field Theory*, Addison-Wesley (1995).
- [10] M.D. Schwartz, *Quantum Field Theory and the Standard Model*, Cambridge University Press (2014).
- [11] F. Halzen and A.D. Martin, *Quarks and leptons: an introductory course in modern particle physics*, Wiley (1984).
- [12] D. Bailin and A. Love, *Introduction to Gauge Field Theory*, Taylor & Francis (1986).
- [13] C. Burgess and G. Moore, *The Standard Model: A primer*, Cambridge University Press (2006).
- [14] J.C. Romao and J.P. Silva, *A resource for signs and Feynman diagrams of the Standard Model*, *International Journal of Modern Physics A* **27** (2012) 1230025 [1209.6213v2].
- [15] S.J. Parke and T. Taylor, *An Amplitude for n Gluon Scattering*, *Phys. Rev. Lett.* **56** (1986) 2459.
- [16] F.A. Berends and W. Giele, *Recursive calculations for processes with n gluons*, *Nucl. Phys. B* **306** (1988) 759.
- [17] H. Elvang and Y. tin Huang, *Scattering Amplitudes in Gauge Theory and Gravity*, Cambridge University Press (2015), [1308.1697v2].
- [18] S. Weinberg, *The Quantum theory of fields. Vol. 1: Foundations*, Cambridge University Press (2005).
- [19] R. Britto, F. Cachazo, B. Feng and E. Witten, *Direct proof of tree-level recursion relation in Yang-Mills theory*, *Phys. Rev. Lett.* **94** (2005) 181602 [hep-th/0501052].

- [20] P.W. Higgs, *Broken symmetries, massless particles and gauge fields*, *Phys. Lett.* **12** (1964) 132.
- [21] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, *Phys. Rev. Lett.* **13** (1964) 321.
- [22] PARTICLE DATA GROUP collaboration, *Review of Particle Physics*, *PTEP* **2020** (2020) 083C01.
- [23] J. Goldstone, *Field Theories with Superconductor Solutions*, *Nuovo Cim.* **19** (1961) 154.
- [24] W. Greiner and B. Müller, *Gauge theory of weak interactions*, Springer (2000).
- [25] N. Cabibbo, *Unitary Symmetry and Leptonic Decays*, *Phys. Rev. Lett.* **10** (1963) 531.
- [26] M. Kobayashi and T. Maskawa, *CP Violation in the Renormalizable Theory of Weak Interaction*, *Prog. Theor. Phys.* **49** (1973) 652.
- [27] CKMFITTER collaboration, *CP violation and the CKM matrix: Assessing the impact of the asymmetric B factories*, *Eur. Phys. J. C* **41** (2005) 1
[hep-ph/0406184]. Updated results and plots available at
<http://ckmfitter.in2p3.fr>.
- [28] UTFIT collaboration, *The Unitarity Triangle Fit in the Standard Model and Hadronic Parameters from Lattice QCD: A Reappraisal after the Measurements of Δm_s and $BR(B \rightarrow \tau \nu_\tau)$* , *JHEP* **10** (2006) 081
[hep-ph/0606167].
- [29] L. Wolfenstein, *Parametrization of the Kobayashi-Maskawa Matrix*, *Phys. Rev. Lett.* **51** (1983) 1945.
- [30] G. Luders, *On the Equivalence of Invariance under Time Reversal and under Particle-Antiparticle Conjugation for Relativistic Field Theories*, *Kong. Dan. Vid. Sel. Mat. Fys. Med.* **28N5** (1954) 1.

-
- [31] J. Bell, *Time reversal in field theory*, *Proc. Roy. Soc. Lond. A* **231** (1955) 479.
- [32] ATLAS collaboration, *Standard Model Summary Plots Spring 2020*, Tech. Rep. ATL-PHYS-PUB-2020-010, CERN, Geneva (May, 2020).
- [33] SUPER-KAMIOKANDE collaboration, *Evidence for oscillation of atmospheric neutrinos*, *Phys. Rev. Lett.* **81** (1998) 1562 [[hep-ex/9807003](#)].
- [34] Y. Grossman, E. Passemar and S. Schacht, *On the Statistical Treatment of the Cabibbo Angle Anomaly*, *JHEP* **07** (2020) 068 [[1911.07821](#)].
- [35] C.-Y. Seng, X. Feng, M. Gorchtein and L.-C. Jin, *Joint lattice QCD–dispersion theory analysis confirms the quark-mixing top-row unitarity deficit*, *Phys. Rev. D* **101** (2020) 111301 [[2003.11264](#)].
- [36] M. Kirk, *Cabibbo anomaly versus electroweak precision tests: An exploration of extensions of the standard model*, *Phys. Rev. D* **103** (2021) 035004 [[2008.03261](#)].
- [37] B. Belfatto, R. Beradze and Z. Berezhiani, *The CKM unitarity problem: A trace of new physics at the TeV scale?*, *Eur. Phys. J. C* **80** (2020) 149 [[1906.02714](#)].
- [38] A. Crivellin, F. Kirk, C.A. Manzari and M. Montull, *Global Electroweak Fit and Vector-Like Leptons in Light of the Cabibbo Angle Anomaly*, *JHEP* **20** (2020) 166 [[2008.01113](#)].
- [39] R.P. Feynman, *Very high-energy collisions of hadrons*, *Phys. Rev. Lett.* **23** (1969) 1415.
- [40] R.K. Ellis, W.J. Stirling and B.R. Webber, *QCD and Collider Physics*, Cambridge University Press (1996).
- [41] J.C. Collins, D.E. Soper and G. Sterman, *Factorization of hard processes in QCD*, [[hep-ph/0409313v1](#)].

- [42] CTEQ collaboration, *Handbook of perturbative QCD: Version 1.0*, *Rev. Mod. Phys.* **67** (1995) 157.
- [43] J.C. Collins and D.E. Soper, *The Theorems of Perturbative QCD*, *Ann. Rev. Nucl. Part. Sci.* **37** (1987) 383.
- [44] J. Collins, *Foundations of perturbative QCD*, vol. 32, Cambridge University Press (2013).
- [45] J.C. Collins, D.E. Soper and G.F. Sterman, *Soft Gluons and Factorization*, *Nucl. Phys. B* **308** (1988) 833.
- [46] M. Stratmann and W. Vogelsang, *Next-to-leading order evolution of polarized and unpolarized fragmentation functions*, *Nucl. Phys. B* **496** (1997) 41 [hep-ph/9612250].
- [47] O. Gituliar and S. Moch, *Towards three-loop QCD corrections to the time-like splitting functions*, *Acta Phys. Polon. B* **46** (2015) 1279 [1505.02901].
- [48] S. Moch, J. Vermaseren and A. Vogt, *The Three loop splitting functions in QCD: The Nonsinglet case*, *Nucl. Phys. B* **688** (2004) 101 [hep-ph/0403192].
- [49] A. Vogt, S. Moch and J. Vermaseren, *The Three-loop splitting functions in QCD: The Singlet case*, *Nucl. Phys. B* **691** (2004) 129 [hep-ph/0404111].
- [50] G. Altarelli and G. Parisi, *Asymptotic Freedom in Parton Language*, *Nucl. Phys. B* **126** (1977) 298.
- [51] H.-W. Lin et al., *Parton distributions and lattice QCD calculations: a community white paper*, *Prog. Part. Nucl. Phys.* **100** (2018) 107 [1711.07916].
- [52] H1, ZEUS collaboration, *Combination of measurements of inclusive deep inelastic $e^\pm p$ scattering cross sections and QCD analysis of HERA data*, *Eur. Phys. J. C* **75** (2015) 580 [1506.06042].

- [53] NNPDF collaboration, *Parton distributions from high-precision collider data*, *Eur. Phys. J. C* **77** (2017) 663 [1706.00428].
- [54] L. Harland-Lang, A. Martin, P. Motylinski and R. Thorne, *Parton distributions in the LHC era: MMHT 2014 PDFs*, *Eur. Phys. J. C* **75** (2015) 204 [1412.3989].
- [55] S. Dulat, T.-J. Hou, J. Gao, M. Guzzi, J. Huston, P. Nadolsky et al., *New parton distribution functions from a global analysis of quantum chromodynamics*, *Phys. Rev. D* **93** (2016) 033006 [1506.07443].
- [56] LHEC, FCC-HE STUDY GROUP collaboration, *The Large Hadron-Electron Collider at the HL-LHC*, [2007.14491].
- [57] T.-J. Hou et al., *New CTEQ global analysis of quantum chromodynamics with high-precision data from the LHC*, *Phys. Rev. D* **103** (2021) 014013 [1912.10053].
- [58] T. Luthe, A. Maier, P. Marquard and Y. Schroder, *Complete renormalization of QCD at five loops*, *JHEP* **03** (2017) 020 [1701.07068].
- [59] P. Baikov, K. Chetyrkin and J. Kühn, *Five-Loop Running of the QCD coupling constant*, *Phys. Rev. Lett.* **118** (2017) 082002 [1606.08659].
- [60] F. Herzog, B. Ruijl, T. Ueda, J. Vermaseren and A. Vogt, *The five-loop beta function of Yang-Mills theory with fermions*, *JHEP* **02** (2017) 090 [1701.01404].
- [61] S. Dittmaier, *Separation of soft and collinear singularities from one loop N point integrals*, *Nucl. Phys. B* **675** (2003) 447 [hep-ph/0308246].
- [62] S. Catani and M.H. Seymour, *A general algorithm for calculating jet cross sections in NLO QCD*, *Nucl. Phys. B* **485** (1997) 291 [hep-ph/9605323v3]. [Erratum: Nucl.Phys.B 510, 503–504 (1998)].

- [63] S. Catani, *The Singular behavior of QCD amplitudes at two loop order*, *Phys. Lett. B* **427** (1998) 161 [[hep-ph/9802439](#)].
- [64] J. Currie, E.W.N. Glover and S. Wells, *Infrared structure at NNLO using antenna subtraction*, *JHEP* **2013** (2013) 66 [[1301.4693v2](#)].
- [65] F. Berends and W. Giele, *Multiple soft gluon radiation in parton processes*, *Nucl. Phys. B* **313** (1989) 595.
- [66] K.J. Ozeren and W.J. Stirling, *MHV techniques for QED processes*, *JHEP* **2005** (2005) 016 [[hep-th/0509063v2](#)].
- [67] J. Campbell and E.W.N. Glover, *Double unresolved approximations to multiparton scattering amplitudes*, *Nucl. Phys. B* **527** (1998) 264 [[hep-ph/9710255v3](#)].
- [68] T. Schuster, *Color ordering in QCD*, *Phys. Rev. D* **89** (2014) 105022 [[1311.6296v1](#)].
- [69] A. Gehrmann-De Ridder, T. Gehrmann and E.W.N. Glover, *Antenna subtraction at NNLO*, *JHEP* **2005** (2005) 056 [[hep-ph/0505111v3](#)].
- [70] A. Gehrmann-De Ridder, T. Gehrmann and M. Ritzmann, *Antenna subtraction at NNLO with hadronic initial states: double real initial-initial configurations*, *JHEP* **2012** (2012) [[1207.5779v1](#)].
- [71] Z. Bern, V. Del Duca, W.B. Kilgore and C.R. Schmidt, *The infrared behavior of one loop QCD amplitudes at next-to-next-to leading order*, *Phys. Rev. D* **60** (1999) 116001 [[hep-ph/9903516](#)].
- [72] G.F. Sterman and S. Weinberg, *Jets from Quantum Chromodynamics*, *Phys. Rev. Lett.* **39** (1977) 1436.
- [73] T. Kinoshita, *Mass singularities of Feynman amplitudes*, *J. Math. Phys.* **3** (1962) 650 [<http://dx.doi.org/10.1063/1.1724268>].

- [74] T.D. Lee and M. Nauenberg, *Degenerate systems and mass singularities*, *Phys. Rev.* **133** (1964) B1549.
- [75] F. Bloch and A. Nordsieck, *Note on the Radiation Field of the electron*, *Phys. Rev.* **52** (1937) 54.
- [76] C. Anastasiou and K. Melnikov, *Higgs boson production at hadron colliders in NNLO QCD*, *Nucl. Phys. B* **646** (2002) 220 [[hep-ph/0207004](#)].
- [77] A. Gehrmann-De Ridder, E.W.N. Glover and J. Pires, *Real-Virtual corrections for gluon scattering at NNLO*, *JHEP* **02** (2012) 141 [[1112.3613](#)].
- [78] M. Cacciari and N. Houdeau, *Meaningful characterisation of perturbative theoretical uncertainties*, *JHEP* **2011** (2011) [[1105.5152v2](#)].
- [79] E. Bagnaschi, M. Cacciari, A. Guffanti and L. Jenniches, *An extensive survey of the estimation of uncertainties from missing higher orders in perturbative calculations*, *JHEP* **02** (2015) 133 [[1409.5036](#)].
- [80] M. Bonvini, *Probabilistic definition of the perturbative theoretical uncertainty from missing higher orders*, *Eur. Phys. J. C* **80** (2020) 989 [[2006.16293](#)].
- [81] P. Jimenez-Delgado, *The role of the input scale in parton distribution analyses*, *Phys. Lett. B* **714** (2012) 301 [[1206.4262](#)].
- [82] L. Harland-Lang and R. Thorne, *On the Consistent Use of Scale Variations in PDF Fits and Predictions*, *Eur. Phys. J. C* **79** (2019) 225 [[1811.08434](#)].
- [83] NNPDF collaboration, *A first determination of parton distributions with theoretical uncertainties*, *Eur. Phys. J. C* (2019) 79:838 [[1905.04311](#)].
- [84] NNPDF collaboration, *Parton Distributions with Theory Uncertainties: General Formalism and First Phenomenological Studies*, *Eur. Phys. J. C* **79** (2019) 931 [[1906.10698](#)].

- [85] G. Grunberg, *Renormalization Group Improved Perturbative QCD*, *Phys. Lett. B* **95** (1980) 70. [Erratum: *Phys.Lett.B* 110, 501 (1982)].
- [86] P.M. Stevenson, *Resolution of the Renormalization Scheme Ambiguity in Perturbative QCD*, *Phys. Lett. B* **100** (1981) 61.
- [87] P.M. Stevenson, *Optimized Perturbation Theory*, *Phys. Rev. D* **23** (1981) 2916.
- [88] P.M. Stevenson and H. Politzer, *Optimized perturbation theory applied to factorization scheme dependence*, *Nucl. Phys. B* **277** (1986) 758.
- [89] P. Stevenson, *Optimization for factorized quantities in perturbative QCD*, *Nucl. Phys. B* **944** (2019) 114635 [1904.07159].
- [90] S.J. Brodsky, G. Lepage and P.B. Mackenzie, *On the Elimination of Scale Ambiguities in Perturbative Quantum Chromodynamics*, *Phys. Rev. D* **28** (1983) 228.
- [91] T. Gehrmann, N. Häfliger and P.F. Monni, *BLM Scale Fixing in Event Shape Distributions*, *Eur. Phys. J. C* **74** (2014) 2896 [1401.6809].
- [92] H.-H. Ma, X.-G. Wu, Y. Ma, S.J. Brodsky and M. Mojaza, *Setting the renormalization scale in perturbative QCD: Comparisons of the principle of maximum conformality with the sequential extended Brodsky-Lepage-Mackenzie approach*, *Phys. Rev. D* **91** (2015) 094028 [1504.01260].
- [93] H.A. Chawdhry and A. Mitov, *Ambiguities of the principle of maximum conformality procedure for hadron collider processes*, *Phys. Rev. D* **100** (2019) 074013 [1907.06610].
- [94] S.J. Brodsky and L. Di Giustino, *Setting the Renormalization Scale in QCD: The Principle of Maximum Conformality*, *Phys. Rev. D* **86** (2012) 085026 [1107.0338].

- [95] S.J. Brodsky and X.-G. Wu, *Scale Setting Using the Extended Renormalization Group and the Principle of Maximum Conformality: the QCD Coupling Constant at Four Loops*, *Phys. Rev. D* **85** (2012) 034038 [1111.6175]. [Erratum: *Phys.Rev.D* 86, 079903 (2012)].
- [96] X.-G. Wu, S.J. Brodsky and M. Mojaza, *The Renormalization Scale-Setting Problem in QCD*, *Prog. Part. Nucl. Phys.* **72** (2013) 44 [1302.0599].
- [97] UA1 collaboration, *Hadronic Jet Production at the CERN Proton - anti-Proton Collider*, *Phys. Lett. B* **132** (1983) 214.
- [98] G.C. Blazey et al., *Run II jet physics*, in *Physics at Run II: QCD and Weak Boson Physics Workshop: Final General Meeting*, pp. 47–77, May, 2000 [hep-ex/0005012].
- [99] G.P. Salam and G. Soyez, *A Practical Seedless Infrared-Safe Cone jet algorithm*, *JHEP* **05** (2007) 086 [0704.0292].
- [100] G.P. Salam, *Towards Jetography*, *Eur. Phys. J. C* **67** (2010) 637 [0906.1833].
- [101] S. Catani, Y.L. Dokshitzer, M. Olsson, G. Turnock and B. Webber, *New clustering algorithm for multi - jet cross-sections in $e^+ e^-$ annihilation*, *Phys. Lett. B* **269** (1991) 432.
- [102] W. Stirling, *Hard QCD working group: Theory summary*, *J. Phys. G* **17** (1991) 1567.
- [103] Y.L. Dokshitzer, G. Leder, S. Moretti and B. Webber, *Better jet clustering algorithms*, *JHEP* **08** (1997) 001 [hep-ph/9707323].
- [104] M. Wobisch and T. Wengler, *Hadronization corrections to jet cross-sections in deep inelastic scattering*, in *Workshop on Monte Carlo Generators for HERA Physics (Plenary Starting Meeting)*, pp. 270–279, Apr., 1998 [hep-ph/9907280].

- [105] M. Cacciari, G.P. Salam and G. Soyez, *The anti- k_t jet clustering algorithm*, *JHEP* **04** (2008) 063 [0802.1189].
- [106] J. Shelton, *Jet Substructure*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Searching for New Physics at Small and Large Scales*, pp. 303–340, 2013, DOI [1302.0260].
- [107] A.J. Larkoski, I. Moult and B. Nachman, *Jet Substructure at the Large Hadron Collider: A Review of Recent Advances in Theory and Machine Learning*, *Phys. Rept.* **841** (2020) 1 [1709.04464].
- [108] R. Kogler et al., *Jet Substructure at the Large Hadron Collider: Experimental Review*, *Rev. Mod. Phys.* **91** (2019) 045003 [1803.06991].
- [109] S. Marzani, G. Soyez and M. Spannowsky, *Looking inside jets: an introduction to jet substructure and boosted-object phenomenology*, vol. 958, Springer (2019), 10.1007/978-3-030-15709-8, [1901.10342].
- [110] A.J. Larkoski, S. Marzani, G. Soyez and J. Thaler, *Soft drop*, *JHEP* **2014** (2014) [1402.2657v2].
- [111] Z. Kunszt, *From scattering amplitudes to cross-sections in QCD*, in *Theoretical Advanced Study Institute in Elementary Particle Physics (TASI 95): QCD and Beyond*, pp. 515–538, Mar., 1996 [hep-ph/9603235].
- [112] W.T. Giele and E.W.N. Glover, *Higher-order corrections to jet cross sections in e^+e^- -annihilation*, *Phys. Rev. D* **46** (1992) 1980.
- [113] W. Giele, E.W.N. Glover and D.A. Kosower, *Higher order corrections to jet cross-sections in hadron colliders*, *Nucl. Phys. B* **403** (1993) 633 [hep-ph/9302225].
- [114] S. Catani and M. Grazzini, *Next-to-next-to-leading-order subtraction formalism in hadron collisions and its application to Higgs-boson production*

- at the large hadron collider*, *Phys. Rev. Lett.* **98** (2007) 222002 [hep-ph/0703012v2].
- [115] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Universality of transverse-momentum resummation and hard factors at the NNLO*, *Nucl. Phys. B* **881** (2014) 414 [1311.1654].
- [116] J.C. Collins, D.E. Soper and G.F. Sterman, *Transverse Momentum Distribution in Drell-Yan Pair and W and Z Boson Production*, *Nucl. Phys. B* **250** (1985) 199.
- [117] S. Catani, D. de Florian and M. Grazzini, *Universality of nonleading logarithmic contributions in transverse momentum distributions*, *Nucl. Phys. B* **596** (2001) 299 [hep-ph/0008184].
- [118] D. de Florian and M. Grazzini, *The Structure of large logarithmic corrections at small transverse momentum in hadronic collisions*, *Nucl. Phys. B* **616** (2001) 247 [hep-ph/0108273].
- [119] T. Becher and M. Neubert, *Drell-Yan Production at Small q_T , Transverse Parton Distributions and the Collinear Anomaly*, *Eur. Phys. J. C* **71** (2011) 1665 [1007.4005].
- [120] I.W. Stewart, F.J. Tackmann and W.J. Waalewijn, *N-Jettiness: An Inclusive Event Shape to Veto Jets*, *Phys. Rev. Lett.* **105** (2010) 092002 [1004.2489].
- [121] M.A. Ebert, I. Moul, I.W. Stewart, F.J. Tackmann, G. Vita and H.X. Zhu, *Power Corrections for N-Jettiness Subtractions at $\mathcal{O}(\alpha_s)$* , *JHEP* **12** (2018) 084 [1807.10764].
- [122] L. Cieri, C. Oleari and M. Rocco, *Higher-order power corrections in a transverse-momentum cut for colour-singlet production at NLO*, *Eur. Phys. J. C* **79** (2019) 852 [1906.09044].

- [123] M.A. Ebert and F.J. Tackmann, *Impact of isolation and fiducial cuts on q_T and N -jettiness subtractions*, *JHEP* **03** (2020) 158 [[1911.08486](#)].
- [124] S. Frixione, Z. Kunszt and A. Signer, *Three jet cross-sections to next-to-leading order*, *Nucl. Phys. B* **467** (1996) 399 [[hep-ph/9512328](#)].
- [125] S. Frixione, *A General approach to jet cross-sections in QCD*, *Nucl. Phys. B* **507** (1997) 295 [[hep-ph/9706545](#)].
- [126] T. Gleisberg and F. Krauss, *Automating dipole subtraction for QCD NLO calculations*, *Eur. Phys. J. C* **53** (2008) 501 [[0709.2881v1](#)].
- [127] R. Frederix, T. Gehrmann and N. Greiner, *Automation of the Dipole Subtraction Method in MadGraph/MadEvent*, *JHEP* **09** (2008) 122 [[0808.2128](#)].
- [128] K. Hasegawa, S. Moch and P. Uwer, *AutoDipole: Automated generation of dipole subtraction terms*, *Comput. Phys. Commun.* **181** (2010) 1802 [[0911.4371](#)].
- [129] R. Frederix, T. Gehrmann and N. Greiner, *Integrated dipoles with MadDipole in the MadGraph framework*, *JHEP* **06** (2010) 086 [[1004.2905](#)].
- [130] S. Platzer and S. Gieseke, *Dipole Showers and Automated NLO Matching in Herwig++*, *Eur. Phys. J. C* **72** (2012) 2187 [[1109.6256](#)].
- [131] E. Bothmann, G.S. Chahal, S. Höche, J. Krause, F. Krauss, S. Kuttimalai et al., *Event generation with Sherpa 2.2*, [[1905.09127v1](#)].
- [132] J. Bellm et al., *Herwig 7.2 release note*, *Eur. Phys. J. C* **80** (2020) 452 [[1912.06509](#)].
- [133] J. Alwall, R. Frederix, S. Frixione, V. Hirschi, F. Maltoni, O. Mattelaer et al., *The automated computation of tree-level and next-to-leading order differential cross sections, and their matching to parton shower simulations*, *JHEP* **07** (2014) 079 [[1405.0301](#)].

- [134] M. Czakon, *A novel subtraction scheme for double-real radiation at NNLO*, *Phys. Lett. B* **693** (2010) 259 [1005.0274].
- [135] M. Czakon and D. Heymes, *Four-dimensional formulation of the sector-improved residue subtraction scheme*, *Nucl. Phys. B* **890** (2014) 152 [1408.2500].
- [136] M. Czakon, A. van Hameren, A. Mitov and R. Poncelet, *Single-jet inclusive rates with exact color at $\mathcal{O}(\alpha_s^4)$* , *JHEP* **10** (2019) 262 [1907.12911].
- [137] A. Behring, M. Czakon and R. Poncelet, *Sector-improved residue subtraction: Improvements and Applications*, *PoS LL2018* (2018) 024 [1808.07656].
- [138] F. Caola, K. Melnikov and R. Röntsch, *Nested soft-collinear subtractions in NNLO QCD computations*, *Eur. Phys. J. C* **77** (2017) 248 [1702.01352].
- [139] Z. Trócsányi, G. Somogyi and F. Tramontano, *Fully Differential Decay Rate of a Standard Model Higgs Boson into Two b-jets at NNLO*, *Acta Phys. Polon. B* **46** (2015) 2097.
- [140] L. Magnea, E. Maina, G. Pelliccioli, C. Signorile-Signorile, P. Torrielli and S. Uccirati, *Local analytic sector subtraction at NNLO*, *JHEP* **12** (2018) 107 [1806.09570]. [Erratum: JHEP 06, 013 (2019)].
- [141] F. Herzog, *Geometric IR subtraction for final state real radiation*, *JHEP* **08** (2018) 006 [1804.07949].
- [142] A. Gehrmann-De Ridder, T. Gehrmann and E.W.N. Glover, *Infrared structure of $e^+e^- \rightarrow 2$ jets at NNLO*, *Nucl. Phys. B* **691** (2004) 195 [hep-ph/0403057].
- [143] A. Gehrmann-De Ridder, T. Gehrmann and E.W.N. Glover, *Quark-gluon antenna functions from neutralino decay*, *Phys. Lett. B* **612** (2005) 36 [hep-ph/0501291].

- [144] A. Gehrmann-De Ridder, T. Gehrmann and E.W.N. Glover, *Gluon-gluon antenna functions from Higgs boson decay*, *Phys. Lett. B* **612** (2005) 49 [hep-ph/0502110].
- [145] C. Duhr, *Mathematical aspects of scattering amplitudes*, in *Theoretical Advanced Study Institute in Elementary Particle Physics: Journeys Through the Precision Frontier: Amplitudes for Colliders*, pp. 419–476, 2015, DOI [1411.7538].
- [146] A. Daleo, T. Gehrmann and D. Maitre, *Antenna subtraction with hadronic initial states*, *JHEP* **04** (2007) 016 [hep-ph/0612257].
- [147] D.A. Kosower, *Multiple singular emission in gauge theories*, *Phys. Rev. D* **67** (2003) 116003 [hep-ph/0212097].
- [148] A. Daleo, A. Gehrmann-De Ridder, T. Gehrmann and G. Luisoni, *Antenna subtraction at NNLO with hadronic initial states: initial-final configurations*, *JHEP* **01** (2010) 118 [0912.0374].
- [149] R. Boughezal, A. Gehrmann-De Ridder and M. Ritzmann, *Antenna subtraction at NNLO with hadronic initial states: double real radiation for initial-initial configurations with two quark flavours*, *JHEP* **02** (2011) 098 [1011.6631].
- [150] T. Gehrmann and P.F. Monni, *Antenna subtraction at NNLO with hadronic initial states: real-virtual initial-initial configurations*, *JHEP* **12** (2011) 049 [1107.4037].
- [151] A. Gehrmann-De Ridder, T. Gehrmann and M. Ritzmann, *Antenna subtraction at NNLO with hadronic initial states: double real initial-initial configurations*, *JHEP* **10** (2012) 047 [1207.5779].
- [152] J.C. Collins, *Renormalization: An Introduction to Renormalization, The Renormalization Group, and the Operator Product Expansion*, vol. 26 of

- Cambridge Monographs on Mathematical Physics*, Cambridge University Press (1986), 10.1017/CBO9780511622656.
- [153] G. 't Hooft and M. Veltman, *Regularization and Renormalization of Gauge Fields*, *Nucl. Phys. B* **44** (1972) 189.
- [154] W. Siegel, *Supersymmetric Dimensional Regularization via Dimensional Reduction*, *Phys. Lett. B* **84** (1979) 193.
- [155] C. Anastasiou, E.W.N. Glover and M. Tejeda-Yeomans, *Two-loop QED and QCD corrections to massless fermion-boson scattering*, *Nucl. Phys. B* **629** (2002) 255 [[hep-ph/0201274v1](#)].
- [156] V.D. Duca, W.B. Kilgore and F. Maltoni, *Multi-photon amplitudes for next-to-leading order QCD*, *Nucl. Phys. B* **574** (2000) 851 [[hep-ph/9910253v1](#)].
- [157] A. Signer, *Helicity method for next-to-leading order corrections in QCD*, Ph.D. thesis, ETH Zürich, 1995.
- [158] J.M. Campbell and C. Williams, *Triphoton production at hadron colliders*, *Phys. Rev. D* **89** (2014) 113001 [[1403.2641](#)].
- [159] D.A. Dicus and S.S.D. Willenbrock, *Photon pair production and the intermediate-mass Higgs boson*, *Phys. Rev. D* **37** (1988) 1801.
- [160] R. Gauld, E.W.N. Glover, A. Huss, I. Majer and A. Gehrmann-De Ridder, *LHC observables with NNLOJET*, *PoS RADCOR2019* (2019) 002.
- [161] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and T. Morgan, *The NNLO QCD corrections to Z boson production at large transverse momentum*, *JHEP* **07** (2016) 133 [[1605.04295](#)].
- [162] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and T. Morgan, *Precise QCD predictions for the production of a Z boson in*

- association with a hadronic jet*, *Phys. Rev. Lett.* **117** (2016) 022001 [1507.02850].
- [163] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and D. Walker, *Next-to-Next-to-Leading-Order QCD Corrections to the Transverse Momentum Distribution of Weak Gauge Bosons*, *Phys. Rev. Lett.* **120** (2018) 122001 [1712.07543].
- [164] A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and D. Walker, *Vector Boson Production in Association with a Jet at Forward Rapidities*, *Eur. Phys. J. C* **79** (2019) 526 [1901.11041].
- [165] X. Chen, T. Gehrmann, E.W.N. Glover, M. Höfer and A. Huss, *Isolated photon and photon+jet production at NNLO QCD accuracy*, *JHEP* **04** (2020) 166 [1904.01044].
- [166] J. Cruz-Martinez, T. Gehrmann, E.W.N. Glover and A. Huss, *Second-order QCD effects in Higgs boson production through vector boson fusion*, *Phys. Lett. B* **781** (2018) 672 [1802.02445].
- [167] X. Chen, T. Gehrmann, E.W.N. Glover and M. Jaquier, *Precise QCD predictions for the production of Higgs + jet final states*, *Phys. Lett. B* **740** (2015) 147 [1408.5325].
- [168] R. Gauld, A. Gehrmann-De Ridder, E.W.N. Glover, A. Huss and I. Majer, *Precise predictions for WH+jet production at the LHC*, [2009.14209].
- [169] J. Currie, E.W.N. Glover and J. Pires, *Next-to-Next-to Leading Order QCD Predictions for Single Jet Inclusive Production at the LHC*, *Phys. Rev. Lett.* **118** (2017) 072002 [1611.01460].
- [170] J. Currie, A. Gehrmann-De Ridder, T. Gehrmann, E.W.N. Glover, A. Huss and J. Pires, *Precise predictions for dijet production at the LHC*, *Phys. Rev. Lett.* **119** (2017) [1705.10271v1].

- [171] T. Gehrmann, E.W.N. Glover, A. Huss, J. Niehues and H. Zhang, *NNLO QCD corrections to event orientation in e^+e^- annihilation*, *Phys. Lett. B* **775** (2017) 185 [1709.01097].
- [172] J. Niehues and D. Walker, *NNLO QCD Corrections to Jet Production in Charged Current Deep Inelastic Scattering*, *Phys. Lett. B* **788** (2019) 243 [1807.02529].
- [173] A. Buckley, J. Ferrando, S. Lloyd, K. Nordström, B. Page, M. Rüfenacht et al., *LHAPDF6: parton density access in the LHC precision era*, *Eur. Phys. J. C* **75** (2015) 132 [1412.7420].
- [174] F. Buccioni, J.-N. Lang, J.M. Lindert, P. Maierhöfer, S. Pozzorini, H. Zhang et al., *OpenLoops 2*, *Eur. Phys. J. C* **79** (2019) 866 [1907.13071].
- [175] C. Berger, Z. Bern, L. Dixon, F. Febres Cordero, D. Forde, H. Ita et al., *An Automated Implementation of On-Shell Methods for One-Loop Amplitudes*, *Phys. Rev. D* **78** (2008) 036003 [0803.4180].
- [176] ATLAS collaboration, *Measurements of Higgs boson properties in the diphoton decay channel with 36 fb^{-1} of pp collision data at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *Phys. Rev. D* **98** (2018) 052005 [1802.04146].
- [177] CMS collaboration, *A measurement of the Higgs boson mass in the diphoton decay channel*, *Phys. Lett. B* **805** (2020) 135425 [2002.06398].
- [178] ATLAS collaboration, *Search for new phenomena in high-mass diphoton final states using 37 fb^{-1} of proton-proton collisions collected at $\sqrt{s} = 13\text{ TeV}$ with the ATLAS detector*, *Phys. Lett. B* **775** (2017) 105 [1707.04147].
- [179] CMS collaboration, *Search for physics beyond the standard model in high-mass diphoton events from proton-proton collisions at $\sqrt{s} = 13\text{ TeV}$* , *Phys. Rev. D* **98** (2018) 092001 [1809.00327].

- [180] ATLAS Collaboration, *Measurements of integrated and differential cross sections for isolated photon pair production in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. D* **95** (2017) [1704.03839v2].
- [181] R.P. Feynman, *Photon-hadron interactions*, Frontiers in physics, CRC Press (1972), 10.1201/9780429493331.
- [182] R. Field and R. Feynman, *Quark Elastic Scattering as a Source of High Transverse Momentum Mesons*, *Phys. Rev. D* **15** (1977) 2590.
- [183] R. Feynman, R. Field and G. Fox, *A Quantum Chromodynamic Approach for the Large Transverse Momentum Production of Particles and Jets*, *Phys. Rev. D* **18** (1978) 3320.
- [184] E.W.N. Glover and A.G. Morgan, *Measuring the photon fragmentation function at LEP*, *Z. Phys. C* **62** (1994) 311.
- [185] A. Mitov, S. Moch and A. Vogt, *Next-to-Next-to-Leading Order Evolution of Non-Singlet Fragmentation Functions*, *Phys. Lett. B* **638** (2006) 61 [hep-ph/0604053].
- [186] S. Moch and A. Vogt, *On third-order timelike splitting functions and top-mediated Higgs decay into hadrons*, *Phys. Lett. B* **659** (2008) 290 [0709.3899].
- [187] A.A. Almasy, S. Moch and A. Vogt, *On the Next-to-Next-to-Leading Order Evolution of Flavour-Singlet Fragmentation Functions*, *Nucl. Phys. B* **854** (2012) 133 [1107.2263].
- [188] ALEPH collaboration, *First measurement of the quark to photon fragmentation function*, *Z. Phys. C* **69** (1996) 365.
- [189] A. Gehrmann-De Ridder and E.W.N. Glover, *A Complete $O(\alpha\alpha_s)$ calculation of the photon + 1 jet rate in e^+e^- annihilation*, *Nucl. Phys. B* **517** (1998) 269 [hep-ph/9707224].

-
- [190] L. Bourhis, M. Fontannaz and J. Guillet, *Quarks and gluon fragmentation functions into photons*, *Eur. Phys. J. C* **2** (1998) 529 [[hep-ph/9704447](#)].
- [191] A. Gehrmann-De Ridder, T. Gehrmann and E. Poulsen, *Measuring the Photon Fragmentation Function at HERA*, *Eur. Phys. J. C* **47** (2006) 395 [[hep-ph/0604030](#)].
- [192] T. Binoth, J. Guillet, E. Pilon and M. Werlen, *A Full next-to-leading order study of direct photon pair production in hadronic collisions*, *Eur. Phys. J. C* **16** (2000) 311 [[hep-ph/9911340](#)].
- [193] E.L. Berger, X.-f. Guo and J.-w. Qiu, *Breakdown of conventional factorization for isolated photon cross-sections*, *Phys. Rev. Lett.* **76** (1996) 2234 [[hep-ph/9512281](#)].
- [194] P. Aurenche, M. Fontannaz, J.P. Guillet, A. Kotikov and E. Pilon, *Is factorization for isolated photon cross sections broken?*, [[hep-ph/9606287v1](#)].
- [195] S. Catani, M. Fontannaz, J.-P. Guillet and E. Pilon, *Cross section of isolated prompt photons in hadron-hadron collisions*, *JHEP* **2002** (2002) 028 [[hep-ph/0204023v3](#)].
- [196] S. Catani, M. Fontannaz and E. Pilon, *Factorization and soft-gluon divergences in isolated-photon cross sections*, *Phys. Rev. D* **58** (1998) [[hep-ph/9803475v2](#)].
- [197] ATLAS collaboration, *Measurement of isolated-photon pair production in pp collisions at $\sqrt{s} = 7$ TeV with the ATLAS detector*, *JHEP* **01** (2013) 086 [[1211.1913](#)].
- [198] S. Frixione, *Isolated photons in perturbative QCD*, *Phys. Lett. B* **429** (1998) 369 [[hep-ph/9801442](#)].

- [199] OPAL collaboration, *Measurement of isolated prompt photon production in photon photon collisions at $\sqrt{s_{ee}} = 183\text{--}209$ GeV*, *Eur. Phys. J. C* **31** (2003) 491 [[hep-ex/0305075](#)].
- [200] SM, NLO MULTILEG WORKING GROUP collaboration, *The SM and NLO Multileg Working Group: Summary report*, in *6th Les Houches Workshop on Physics at TeV Colliders*, pp. 21–189, Mar., 2010 [[1003.1241](#)].
- [201] Z. Hall and J. Thaler, *Photon isolation and jet substructure*, *JHEP* **09** (2018) 164 [[1805.11622](#)].
- [202] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Diphoton production at hadron colliders: a fully-differential QCD calculation at NNLO*, *Phys. Rev. Lett.* **108** (2012) 072001 [[1110.2375](#)]. [Erratum: *Phys.Rev.Lett.* **117**, 089901 (2016)].
- [203] J.M. Campbell, R. Ellis, Y. Li and C. Williams, *Predictions for diphoton production at the LHC through NNLO in QCD*, *JHEP* **07** (2016) 148 [[1603.02663](#)].
- [204] H.A. Chawdhry, M.L. Czakon, A. Mitov and R. Poncelet, *NNLO QCD corrections to three-photon production at the LHC*, *JHEP* **02** (2020) 057 [[1911.00479](#)].
- [205] CMS Collaboration, *Measurement of differential cross sections for the production of a pair of isolated photons in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **74** (2014) [[1405.7225v2](#)].
- [206] F. Siegert, *A practical guide to event generation for prompt photon production with Sherpa*, *J. Phys. G* **44** (2017) 044007 [[1611.07226](#)].
- [207] S. Catani, L. Cieri, D. de Florian, G. Ferrera and M. Grazzini, *Diphoton production at the LHC: a QCD study up to NNLO*, *JHEP* **04** (2018) 142 [[1802.02095](#)].

- [208] S. Frixione and G. Ridolfi, *Jet photoproduction at HERA*, *Nucl. Phys. B* **507** (1997) 315 [[hep-ph/9707345](#)].
- [209] S. Catani and B. Webber, *Infrared safe but infinite: Soft gluon divergences inside the physical region*, *JHEP* **10** (1997) 005 [[hep-ph/9710333](#)].
- [210] M. Grazzini, S. Kallweit and M. Wiesemann, *Fully differential NNLO computations with MATRIX*, *Eur. Phys. J. C* **78** (2018) 537 [[1711.06631](#)].
- [211] A. Banfi and M. Dasgupta, *Dijet rates with symmetric $E(t)$ cuts*, *JHEP* **01** (2004) 027 [[hep-ph/0312108](#)].
- [212] ATLAS collaboration, *Measurements of integrated and differential cross sections for isolated photon pair production in pp collisions at $\sqrt{s} = 8$ TeV with the ATLAS detector*, *Phys. Rev. D* **95** (2017) 112005 [[1704.03839](#)].
- [213] CDF collaboration, *Measurement of the Cross Section for Prompt Isolated Diphoton Production Using the Full CDF Run II Data Sample*, *Phys. Rev. Lett.* **110** (2013) 101801 [[1212.4204](#)].
- [214] D0 collaboration, *Measurement of the differential cross sections for isolated direct photon pair production in $p\bar{p}$ collisions at $\sqrt{s} = 1.96$ TeV*, *Phys. Lett. B* **725** (2013) 6 [[1301.4536](#)].
- [215] CMS collaboration, *Measurement of differential cross sections for the production of a pair of isolated photons in pp collisions at $\sqrt{s} = 7$ TeV*, *Eur. Phys. J. C* **74** (2014) 3129 [[1405.7225](#)].
- [216] J.M. Campbell, R.K. Ellis and C. Williams, *Direct Photon Production at Next-to-Next-to-Leading Order*, *Phys. Rev. Lett.* **118** (2017) 222001 [[1612.04333](#)]. [Erratum: *Phys.Rev.Lett.* 124, 259901 (2020)].
- [217] J.M. Campbell, J. Rojo, E. Slade and C. Williams, *Direct photon production and PDF fits reloaded*, *Eur. Phys. J. C* **78** (2018) 470 [[1802.03021](#)].

- [218] R. Abdul Khalek et al., *Phenomenology of NNLO jet production at the LHC and its impact on parton distributions*, *Eur. Phys. J. C* **80** (2020) 797 [2005.11327].
- [219] S. Badger, A. Guffanti and V. Yundin, *Next-to-leading order QCD corrections to di-photon production in association with up to three jets at the Large Hadron Collider*, *JHEP* **03** (2014) 122 [1312.5927].
- [220] S. Kallweit, J.M. Lindert, P. Maierhöfer, S. Pozzorini and M. Schönherr, *NLO electroweak automation and precise predictions for W +multijet production at the LHC*, *JHEP* **04** (2015) 012 [1412.5157].
- [221] M. Chiesa, N. Greiner and F. Tramontano, *Automation of electroweak corrections for LHC processes*, *J. Phys. G* **43** (2016) 013002 [1507.08579].
- [222] A. Denner and S. Pozzorini, *One-loop leading logarithms in electroweak radiative corrections: I. results*, *Eur. Phys. J. C* **18** (2001) 461 [hep-ph/0104127v1].
- [223] A. Denner and S. Pozzorini, *One-loop leading logarithms in electroweak radiative corrections: II. factorization of collinear singularities*, *Eur. Phys. J. C* **21** (2001) 63 [hep-ph/0104127v1].
- [224] E. Bothmann and D. Napoletano, *Automated evaluation of electroweak Sudakov logarithms in Sherpa*, *Eur. Phys. J. C* **80** (2020) 1024 [2006.14635].
- [225] M. Chiesa, N. Greiner, M. Schönherr and F. Tramontano, *Electroweak corrections to diphoton plus jets*, *JHEP* **10** (2017) 181 [1706.09022].
- [226] G. Cullen, N. Greiner, G. Heinrich, G. Luisoni, P. Mastrolia, G. Ossola et al., *Automated One-Loop Calculations with GoSam*, *Eur. Phys. J. C* **72** (2012) 1889 [1111.2034].

- [227] G. Cullen et al., *GOSAM-2.0: a tool for automated one-loop calculations within the Standard Model and beyond*, *Eur. Phys. J. C* **74** (2014) 3001 [1404.7096].
- [228] M. Höfer, *Higher Order Corrections to Photon Production in Hadronic Collisions at the LHC*, Ph.D. thesis, University of Zurich, 2020.
- [229] D. de Florian, G.F. Sborlini and G. Rodrigo, *QED corrections to the Altarelli–Parisi splitting functions*, *Eur. Phys. J. C* **76** (2016) 282 [1512.00612].
- [230] A. Martin, R. Roberts, W. Stirling and R. Thorne, *Parton distributions incorporating QED contributions*, *Eur. Phys. J. C* **39** (2005) 155 [hep-ph/0411040].
- [231] A. Martin and M. Ryskin, *The photon PDF of the proton*, *Eur. Phys. J. C* **74** (2014) 3040 [1406.2118].
- [232] L. Harland-Lang, A. Martin, R. Nathvani and R. Thorne, *Ad Lucem: QED Parton Distribution Functions in the MMHT Framework*, *Eur. Phys. J. C* **79** (2019) 811 [1907.02750].
- [233] NNPDF collaboration, *Parton distributions with QED corrections*, *Nucl. Phys. B* **877** (2013) 290 [1308.0598].
- [234] NNPDF collaboration, *Illuminating the photon content of the proton within a global PDF analysis*, *SciPost Phys.* **5** (2018) 008 [1712.07053].
- [235] A. Manohar, P. Nason, G.P. Salam and G. Zanderighi, *How bright is the proton? A precise determination of the photon parton distribution function*, *Phys. Rev. Lett.* **117** (2016) 242002 [1607.04266].
- [236] A.V. Manohar, P. Nason, G.P. Salam and G. Zanderighi, *The Photon Content of the Proton*, *JHEP* **12** (2017) 046 [1708.01256].

- [237] C. Schmidt, J. Pumplin, D. Stump and C. Yuan, *CT14QED parton distribution functions from isolated photon production in deep inelastic scattering*, *Phys. Rev. D* **93** (2016) 114015 [1509.02905].
- [238] A. Denner, S. Dittmaier, M. Roth and D. Wackeroth, *Predictions for all processes $e^+e^- \rightarrow 4 \text{ fermions} + \gamma$* , *Nucl. Phys. B* **560** (1999) 33 [hep-ph/9904472].
- [239] A. Denner, S. Dittmaier, M. Roth and L. Wieders, *Electroweak corrections to charged-current $e^+e^- \rightarrow 4 \text{ fermion}$ processes: Technical details and further results*, *Nucl. Phys. B* **724** (2005) 247 [hep-ph/0505042]. [Erratum: *Nucl.Phys.B* 854, 504–507 (2012)].
- [240] A. Banfi, S. Redford, M. Vesterinen, P. Waller and T. Wyatt, *Optimisation of variables for studying dilepton transverse momentum distributions at hadron colliders*, *Eur. Phys. J. C* **71** (2011) 1600 [1009.1580].
- [241] J.C. Collins and D.E. Soper, *Angular Distribution of Dileptons in High-Energy Hadron Collisions*, *Phys. Rev. D* **16** (1977) 2219.
- [242] V. Hirschi, R. Frederix, S. Frixione, M.V. Garzelli, F. Maltoni and R. Pittau, *Automation of one-loop QCD corrections*, [1103.0621v2].
- [243] R. Kleiss, W. Stirling and S. Ellis, *A New Monte Carlo Treatment of Multiparticle Phase Space at High-energies*, *Comput. Phys. Commun.* **40** (1986) 359.
- [244] M.A. Ebert and F.J. Tackmann, *Impact of isolation and fiducial cuts on q_T and N -jettiness subtractions*, *JHEP* **03** (2020) 158 [1911.08486v1].