

Measurement of the photon polarisation in $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays at LHCb

Thèse n. 10426 2025
présentée le 16 mai 2025
à la Faculté des sciences de base
Laboratoire de physique des hautes energies OS
Programme doctoral en Physique
École polytechnique fédérale de Lausanne
pour l'obtention du grade de Docteur ès Sciences
par

Marie Bachmayer

acceptée sur proposition du jury :

Prof. Christian Theiler, président du jury
Prof. Olivier Schneider, directeur de thèse
Prof. Martino Borsato, rapporteur
Dr. Olivier Deschamps, rapporteur
Prof. Lea Caminada, rapporteur

Lausanne, EPFL, 2025

EPFL

Nothing that actually occurs is of the smallest importance.

—Oscar Wilde

Acknowledgements

Explaining how I ended up here is like making an apple pie from scratch. Even limiting myself to my personal experiences, I could go into subatomic detail. While I was not wrong to expect that studying particle physics would help me understand how the world works at a fundamental level, understanding how I got here requires a different tool. Maybe it is cheating to use hindsight to tell a story, but it will allow me to pick out only a small handful of the turning points that stand out to me as I look back, leading to a necessarily non-exhaustive list of the people and – as it will turn out – the places that have led me to the moment where I write these words.

My early interest in physics was sparked not in school but by my dad, who was the first person to intrigue me with stories about wave-particle duality. The question of the subject of my graduate studies was settled when, as a teenager, I was given a tour of CERN together with family and friends by Horst Wenninger. Horst introduced me to an unnamed senior physicist, whose rather careless dismissal of my newly formed ambition to study physics did nothing to deter me. More than anything else, it was the atmosphere of the place, the spirit of collaboration rather than competition, the blackboards full of formulae I glimpsed in the offices of theorists who had left their doors ajar. I would not return to CERN until years later, after Prof. Günther Dissertori's lecture on nuclear and particle physics reminded me why I chose to study physics in the first place. Many of my colleagues disagree with me on the aesthetic merit of the CERN campuses: I am always stalking the grounds with my phone camera in front of my face, trying to find the right framing for the shot that will show the beauty I see. This perhaps has more to do with the fact that CERN has kept the aura of magic and promise I perceived on my first visit there than with the architecture of the lab and office buildings. It was Francesco Polci who got me started at LHCb, and I stayed in the collaboration for my PhD. Any narrative structure will now be abandoned as I continue to express my gratitude to those who have helped, supported, and enabled me by putting their names down on paper in no particular order.

My first shout out goes to the bosons: David Adame Carillo, Sjoerd Bouma, Alberto Giampaolo, Giulia Isabella, and Jonas Kunath. Bob Knighton and Gitanjali Poddar, too: it has been a great joy and an honour to walk into and out of the PhD journey with you. Kaustuv Datta, your support means so much to me. Carina Trippl and Surapat Ek-In, the coolest office mates I could have asked for. Anna Mascellani, Sonia Bouchiba, Sebastian Schulte, Vlado Macko, Ettore Zaffaroni, Yunxuan Song, Raphael, Pavol, Veronica, Gianluca, Jennifer, Lino, Serhii, Elena, Vitalii, I cannot name the whole of the fellowship of PhD students and post-docs at

Acknowledgements

LPHE who brighten my days, but I mean all of you. I have been fortunate to find many friends in other fields at EPFL, too: Annie Guillaume in particular helped me preserve some sanity during the earliest, most confusing stages of the Covid lockdowns, and it is ultimately thanks to her that I met Lukas Huber, Sho Watanabe, Petros Liakopoulos, Moahan Murugappan, and others whom I've had fun adventures. Mrudhula Baskaran, Ziga Casar, Yuji Takabayashi: thank you for countless dinners and workouts, truly you are the best of the best. Friends old and new who I am grateful for and will continue to rely on: Nerea Borreguero Muñoz, Ivana Cvijanović, Shant Dakessian, Zsófia Hesketh, Darren Kelly, Anja Logo, Gautam Maitra, Luka Mališa, Luka Manola, Vanesa Rocha Martin, Steve Matsumoto, An-phi Nguyen and the rest of the Threads, Lydia Paulos, Josefine Quack, Raphaëlle Sabouraud and the rest of CPAL, Mahsa Sanati, Miriam Schmatolla, Shantam Taneja, Matteo Turchetta, Tony and family. Thank you does not even cover the half of it for Andy, who made me dinner and bore with me when my head was stuck in my laptop. My friends at UZH, thanks for bringing me unofficially but unreservedly into the fold, and for all the good times. I am honoured to have been a part of the LHCb collaboration and especially to have got to work with Alison Tully and Arvind Venkateswaran on the magnetic field map. None of this would have been possible without Violaine Bellée, Preema Pais, and María Vieites Díaz. I am endlessly lucky, honoured and pleased to be a part of this team. I similarly cannot write down all the ways I am grateful to and for Prof. Olivier Schneider, who is infinitely reliable and supportive, always has time for me no matter how much else is on his plate, and who made this thesis so much more coherent, complete, and correct. Any imperfections are of course entirely on me. Prof. Fred Blanc's encouragement always makes me believe that any problem can be solved. Prof. Lesya Shchutska gave me the best tips. Prof. Radoslav Marchevski is interested in any discussion. Esther Hofmann and Corinne Craman can invariably sort anything out: they are simply invaluable, and so nice about it, too! Thanks to my thesis jury for their excellent questions, in comparison to which my answers pale, as well they should. Thanks to the Diepers family in Aarau for having my back. Danke Mama, Papa, Vinz, Sevi: weil es euch gibt, bin ich nie wirklich einsam. Finally, thanks to everyone who used to ask me when I would be finishing the PhD for motivating me to make them stop asking. Many of you have now moved on to the question of what I am going to do next— I am working on that, too.

Zürich, June 11, 2025

M. B.

Abstract

All physical processes can be described, at a fundamental level, through interactions of elementary particles. The Standard Model of particle physics (SM) describes three of the four known fundamental forces and all of the known elementary particles. The accuracy of the predictions made using this theory has been experimentally established over a wide range of processes. Physicists endeavour to extend the SM to cover those phenomena, observed or hypothesised, that are currently beyond its scope. To develop these extensions, it is useful to identify processes involving SM particles where SM predictions differ from what is measured in nature.

The fact that flavour-changing neutral currents (FCNC) processes are forbidden at tree level in the SM makes them a promising laboratory for precision measurements that can show effects of physics beyond the SM: New Physics (NP) at energy scales higher than those of an FCNC process can nevertheless indirectly affect it to a potentially measurable extent. This work realises one such precision measurement. It concerns the polarisation of the photon in $b \rightarrow s\gamma$ transitions. Due to the chiral nature of the weak interaction, which is the only SM interaction that mediates FCNC processes, the photon is almost entirely left-handed in decays that proceed through a $b \rightarrow s\gamma$ transition. NP with a chiral structure that is different from that of the weak interaction could enhance the right-handed component of the photon. The photon polarisation in $b \rightarrow s\gamma$ transitions is experimentally constrained through measurements of the inclusive branching fraction $B \rightarrow X_s\gamma$ and other measurements of radiative decays of b hadrons, as well as from $B^0 \rightarrow K^{*0}e^+e^-$ decays at low q^2 . It has been directly measured in $\Lambda_b \rightarrow \Lambda\gamma$ decays. The experimental results are consistent with SM predictions. A measurement of the up-down asymmetry in $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decays has shown that the photon in these decays is polarised, but the value of the polarisation parameter λ_γ could not be extracted from this measurement. In this work, an amplitude analysis of $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decays recorded at the LHCb experiment during Run 1 and Run 2 of the LHC is performed, and a model of the $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decay, including the value of λ_γ , is defined. This measurement further constrains the photon polarisation in $b \rightarrow s\gamma$ transitions and thus contributes to the identification of promising avenues for extending the SM.

The reconstruction of charged particle momenta at LHCb requires a precise knowledge of the magnetic field that bends the particles' trajectories. A new map of the magnetic field of the LHCb dipole was developed for use in Run 3 of the LHC. This map and its development is also discussed in this work.

Keywords: particle physics, LHCb, LHC, flavour physics, FCNC, radiative decays, photon polarisation, amplitude analysis, magnetic field map

Résumé

Tous les processus physiques peuvent être décrits, au niveau fondamental, comme des interactions entre particules élémentaires. Le modèle standard de la physique des particules (MS) inclut trois des quatre forces fondamentales connues, et toutes les particules élémentaires connues. Il prédit les probabilités des interactions entre particules avec une grande précision et a été vérifié expérimentalement dans de nombreuses situations. Cependant, il reste quelques lacunes dans la description des processus naturels qui nécessitent l'extension du modèle standard. Des telles extensions pourraient également décrire des phénomènes non observés à ce jour. L'identification de processus dont les prédictions du modèle standard deviennent des mesures expérimentales est essentiel pour le développement de ces extensions potentielles.

Les courants neutres changeant la saveur (FCNC, flavour changing neutral current) nécessitent, dans le MS, une boucle de particules virtuelles. Dans une telle boucle peuvent contribuer de nouvelles particules ou interactions (hors du MS) dont les effets directs ne deviennent observables qu'à des énergies beaucoup plus importantes que celles du processus FCNC concerné. Ceci fait des processus FCNC un outil de choix pour la recherche de phénomènes au-delà du MS. Cette thèse concerne la polarisation des photons dans les transitions $b \rightarrow s\gamma$, qui sont dues à un courant FCNC. L'interaction faible est la seule interaction du MS à autoriser les FCNC. Sa structure chirale fait que les photons sont presque toujours gauchers dans les transitions $b \rightarrow s\gamma$. Des particules ou interactions d'une structure chirale différente de celle de l'interaction faible peuvent diluer la polarisation des photons. Le paramètre de polarisation λ_γ a déjà été restreint expérimentalement par des mesures du rapport d'embranchement des désintégrations $B \rightarrow X_s\gamma$ et autres désintégrations radiatives de hadrons b , ainsi que dans la désintégration $B^0 \rightarrow K^{*0}e^+e^-$ à petit q^2 . Une mesure directe de la polarisation du photon a été effectuée en utilisant la désintégration $\Lambda_b \rightarrow \Lambda\gamma$. Les résultats experimentaux concordent avec les prédictions du MS. Une étude de l'asymétrie du photon dans la désintégration $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ a montré que λ_γ est non nul, c'est-à-dire que le photon est polarisé. Cependant, cette étude n'a pas permis de déterminer la valeur de λ_γ . Cette thèse, basée sur les données enregistrées par l'expérience LHCb au cours des deux premières périodes d'exploitation du LHC, présente une analyse en amplitudes de la désintégration $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ qui permet une mesure de λ_γ . Cette analyse restreint encore mieux la polarisation des photons dans les transitions $b \rightarrow s\gamma$ et contribue ainsi à l'identification d'avenues prometteuses pour étendre le modèle standard.

La reconstruction des quantités de mouvement des particules chargées à LHCb nécessite une connaissance précise du champ magnétique de l'aimant dipolaire de l'expérience. Le développement d'une nouvelle carte du champ magnétique pour la troisième période d'exploitation du LHC est également présenté dans cette thèse.

Mots-clés : physique des particules, LHCb, LHC, physique des saveurs, FCNC, désintégrations radiatives, polarisation du photon, analyse d'amplitudes, carte du champ magnétique

Zusammenfassung

Allen physikalischen Vorgängen liegen Wechselwirkungen von Elementarteilchen zugrunde. Das Standardmodell der Teilchenphysik (SM) umfasst drei der vier bekannten fundamentalen Wechselwirkungen und alle bekannten Elementarteilchen. Die Übereinstimmung der Vorhersagen des SM zum Ablauf von physikalischen Vorgängen mit Messergebnissen ist in verschiedensten Prozessen fast flächendeckend bestätigt. Um die verbleibenden Lücken in der Beschreibung von Naturprozessen zu schliessen, oder um bislang unbeobachtete Phänomene, die ausserhalb des Geltungsbereichs des SM liegen, zu beschreiben, suchen Physiker nach Erweiterungen des SM. Das Auffinden und die genaue Charakterisierung von Vorgängen, bei denen Messergebnisse von SM-Vorhersagen abweichen, trägt wesentlich zur Entwicklung solcher Erweiterungen bei.

FCNC (*flavour-changing neutral current*)-Prozesse sind vielversprechende Kandidaten für solch abweichende Vorgänge, da zu ihrer Beschreibung im SM Schleifen virtueller Teilchen nötig sind. Neue Physik (Elementarteilchen oder Wechselwirkungen, die nicht im SM enthalten sind), die erst bei Energieskalen direkt beobachtbar wird, die deutlich über der des relevanten FCNC-Prozesses liegen, kann möglicherweise trotzdem in diesen Schleifen auftauchen. In dieser Arbeit wird eine solche Präzisionsmessung durchgeführt. Sie betrifft die Polarisierung des Photons in $b \rightarrow s\gamma$ -Übergängen. FCNC-Prozesse können im SM nur durch die schwache Wechselwirkung stattfinden, welche die Paritätserhaltung maximal verletzt. Daher sind Photonen in $b \rightarrow s\gamma$ -Übergängen fast ausschliesslich linkshändig polarisiert. Neue Physik mit einer Chiralitätsstruktur, die von der schwachen Wechselwirkung abweicht, könnte den Anteil an rechtshändigen Photonen steigern. Messungen der Zerfallsrate von $B \rightarrow X_s\gamma$ - und anderen Strahlungszerfällen von b -Hadronen sowie von $B^0 \rightarrow K^{*0}e^+e^-$ -Zerfällen mit niedrigem q^2 beschränken den Wert des Polarisierungsparameters λ_γ . In $\Lambda_b \rightarrow \Lambda\gamma$ -Zerfällen wurde er direkt gemessen. Die Messresultate sind mit der SM-Vorhersage vereinbar. Die Asymmetrie von Photonen in $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ -Zerfällen zeigt ihre Polarisierung, λ_γ ist also ungleich null. Um den Wert von λ_γ in solchen Zerfällen zu messen ist jedoch eine Amplitudenanalyse notwendig, die in dieser Arbeit durchgeführt wird. Dazu werden Daten untersucht, die in den ersten beiden Runs des LHC am LHCb-Experiment aufgezeichnet wurden. Diese Messung schränkt die Polarisierung des Photons in $b \rightarrow s\gamma$ -Übergängen weiter ein und trägt so zur Auffindung von Hinweisen zur Erweiterung des SM bei.

Das Magnetfeld des LHCb-Dipolmagneten, welches die Bahnen geladener Teilchen krümmt, muss präzise bestimmt sein, damit die Teilchenimpulse genau gemessen werden können. Eine neue Karte dieses Magnetfelds für Run 3 des LHC wurde im Zuge dieser Arbeit entwickelt.

Stichwörter: Teilchenphysik, LHCb, LHC, Flavour-Physik, FCNC, Strahlungszerfälle, Photonpolarisierung, Amplitudenanalyse, Magnetfeldkarte

Contents

Acknowledgements	i
Abstract (English/Français/Deutsch)	iii
1 Introduction	1
2 Theory	5
2.1 The Standard Model of particle physics	5
2.2 S-matrix elements in a hurry	8
2.3 The photon polarisation in $B \rightarrow K\pi\pi\gamma$ decays	10
3 The LHCb spectrometer	15
4 The LHCb magnetic field map	25
4.1 Overview of LHCb magnetic field maps	27
4.2 January 2021 magnetic field map measurements	28
4.3 Simulations of the magnetic field	35
4.3.1 Changes to the simulation model based on surveys in the cavern	36
4.3.2 Additional ferromagnetic structures in the simulation model	37
4.4 Interpolation and extrapolation	39
4.5 The magnetic field map for Run 3	41
5 Data selection and B mass fit	47
5.1 Selection requirements	47
5.2 Fit to B candidate invariant mass	56
5.3 Background subtraction with custom orthogonal weight functions	65
6 Amplitude analysis	69
6.1 Fit function	71
6.2 Fit fractions, visualisation and goodness of fit	75
6.3 Data and MC samples	78
6.4 Iterative amplitude selection	81
6.5 Validation of the amplitude fit method	84
6.6 Results	88
6.7 Cross-checks	94

Contents

6.8 Systematic uncertainties	99
6.8.1 Momentum resolution	100
6.8.2 Fixed parameters in the B mass fit	100
6.8.3 Contamination from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events	102
6.8.4 Integration sample size	110
6.8.5 Choice of decay amplitudes	110
6.8.6 The parametrisation of the $\rho^0 - \omega^0$ resonance mixing	110
6.8.7 Fixed parameters in the amplitude fit	111
7 Conclusion	113
A Appendix	117
A.1 Magnetic field visualisations	118
A.2 Propagators	124
A.2.1 Breit-Wigner	124
A.2.2 Gounaris-Sakurai mass propagator	125
A.2.3 A combined propagator for ρ and ω	125
A.2.4 The LASS line-shape	125
A.2.5 The Bugg line-shape	126
A.3 Elements of the strong decay spin factors	127
Bibliography	129
Curriculum Vitae	135

Chapter 1

Introduction

In 1844, Austrian mineralogist Wilhelm Haidinger was contemplating plates of anadalusite crystals when he observed “a flying phantom of a yellowish color” that vanished again when he tried to focus on it [1] (translation from Ref. [2]). Haidinger continued to observe this phenomenon in reflected light from other materials, and in the refracted light of the blue day-time sky, and correctly concluded that his vision was sensitive to the linear polarisation of light. The image first described by Haidinger is now known as “Haidinger’s brushes” and consists of two perpendicularly intersecting bow-tie shaped bars of yellow and purplish-blue. A representation of this entoptic phenomenon (like the floaters one can see against bright monochromatic backgrounds, the origin of Haidinger’s brushes lies inside the eye itself, so they cannot be photographed) is depicted in Figure 1.1. A seeing person who wishes to test their own sensitivity to light polarisation may look at a white LCD screen and tilt their head, or the screen, from side to side while alert for the appearance of a faint yellow and blue Maltese cross-like structure. The tilting motion rotates the cross with respect to the eye, as the yellow bar always remains perpendicular to the direction of polarisation of the light from the screen, and thus mitigates the neural adaptation that would cause a static image to fade.

This is very cool, but what is it for? It is possible in principle to use Haidinger’s brushes for navigation, to infer the position of the sun [4]: the yellow bars of the cross seen at two different points in the sky lie on celestial great circles that intersect in the sun. For this to be useful,



Figure 1.1 – Haidinger’s brushes [3]. When this page is printed on A4 paper and held at arm’s length, the image is about the same size as the entoptic phenomenon.

however, the sun must not be visible, but the sky nonetheless light enough and its degree of polarisation high enough to make out the faint brushes. If Haidinger's brushes are not very useful, though, why are they cool? Even if there is no direct, tangible benefit to knowing in which direction the light is polarised, the knowledge itself is satisfying. It also raises questions on the mechanism of the perception of Haidinger's brushes, and on the origin and/or reason for the polarisation of the light. Several sections of this work concern a question that sounds similar to the latter, but before posing it, let us take a step back and set the scene. Or rather, if the scene is the physical world, let us break it down into ever smaller and more fundamental parts and processes, until we reach elementary particles and their interactions.

The Standard Model (SM) of particle physics describes the interactions of all known elementary particles, with the exception of gravity. The gravitational force is often compared to the Coulomb force due to their similar mathematical structure when electric charges take the place of masses, but while the carrier particle of the Coulomb force is identified in the SM with the photon, no experimental evidence has so far shown the existence of its gravitational analogue, the graviton, and gravity therefore remains disconnected from the SM. Likewise outside of the scope of the theory are observed phenomena such as neutrino masses, dark matter, dark energy, and matter-antimatter asymmetry (the observable universe is mostly matter and not antimatter, but both should have been produced in equal amounts in the Big Bang). Internally, however, SM predictions agree with experimental measurements across the board. It is therefore natural to want to extend this resoundingly successful theory to incorporate the abovementioned phenomena. Such additions to the SM can take the form of new elementary particles and/or interactions. These must have some connection to the known SM particles if they are to affect the observable universe, and if nature is not fully described by the SM, some deviations from SM predictions ought to arise from the new physics. Correspondingly, if a measurement differs from its SM prediction, it offers a clue to the nature of the necessary extension of the SM. Precision measurements of SM parameters present a way to suss out the weak points of the theory. One such parameter is the photon polarisation parameter in the transition of a bottom quark to a strange quark and a photon (charge-parity conjugation is implied throughout this document except where stated otherwise). Photons can have one of two polarisation states, but due to a quirk of the SM which will be elaborated in Chapter 2, this transition almost exclusively produces photons of the same single polarisation. Human vision cannot detect this polarisation for many reasons, the least of which is that it is in the circular and not the linear basis, so the measurements must be made at particle physics experiments. The photon polarisation in $b \rightarrow s\gamma$ decays has been constrained through measurements of mixing-induced CP asymmetries in radiative decays of B^0 and B_s^0 mesons [5–8] and through the angular analysis of $B^0 \rightarrow K^{*0} e^+ e^-$ decays at invariant masses of the electron pair low enough for the decay to be dominated by an intermediate virtual photon [9].

In 2014, the LHCb collaboration reported the first observation of the photon polarisation in $b \rightarrow s\gamma$ transitions using an analysis of $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays [10]. The amplitude analysis discussed in this work concerns the same decay channel. The 2014 measurement found a non-zero up-down asymmetry of the photon direction with respect to the hadronic decay

plane in the $K\pi\pi$ rest frame. This asymmetry is proportional to the photon polarisation parameter λ_γ [11] (see also Chapter 2), but the determination of the proportionality coefficient requires a comprehensive model of the hadronic part of the decay. This decay mode was studied with a Dalitz analysis to obtain information on the content of the hadronic system and the relative importance of the different decay channels [12]. A method to model the full decay including the angular distributions of the final-state particles in order to measure the photon polarisation parameter by means of an amplitude analysis was proposed in 2019 [13]. The LHCb collaboration reported the strongest direct constraint of the photon polarisation parameter α_γ from an angular analysis of $\Lambda_b^0 \rightarrow \Lambda^0 (\rightarrow p\pi)\gamma$ decays in 2022 [14]. As far as the theoretical description of the quark transition $b \rightarrow s\gamma$ is concerned, the definition of α_γ is the same as that of the parameter λ_γ used in this work, but the two parameters differ in the hadronic corrections required by and the uncertainties associated with their respective hadron decays. When hadronic corrections are neglected, the value predicted at leading order by the SM for both α_γ and λ_γ is $(1 - |r|^2)/(1 + |r|^2)$ where r is the ratio of the quark masses m_s/m_b . The measurement in Ref. [14] of $\alpha_\gamma = 0.82^{+0.17+0.04}_{-0.26-0.13}$, where the first uncertainty is statistical and the second systematic, is consistent with the SM expectation. The amplitude analysis of $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decays proposed in Ref. [13] is sensitive to λ_γ with an expected statistical uncertainty of 0.014 (0.009) for a sample of 14 000 (70 000) $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ events. The implementation of this method discussed in this work is based on a sample of about 35 000 events from data recorded by the LHCb experiment in 2011, 2012, 2015, 2016, 2017, and 2018, so the statistical uncertainty on λ_γ in this measurement is expected to be significantly lower than that on α_γ in the measurement in Ref. [14].

This document is structured as follows: In Chapter 2, the theory of the amplitude analysis of $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decays is motivated. Chapter 3 describes the LHCb detector. Chapter 4 is a detour through the magnetic field map of the LHCb experiment. Chapter 5 returns to the amplitude analysis by describing the selection and B mass fit of the data samples. The amplitude analysis is detailed and its results are discussed in Chapter 6. Chapter 7 concludes the thesis.

Chapter 2

Theory

2.1 The Standard Model of particle physics

In the first year of our education towards a degree in physics, we learn how to derive the equations of motion for a physical system governed by Newtonian mechanics from Newton's laws. Then, in the second or third year, we are introduced to an alternative method of deriving the same equations of motion that is compelling in its elegance: the Lagrangian method. The Lagrangian, named after Italian-French mathematician and physicist Joseph-Louis Lagrange, is defined as

$$L = T - V, \quad (2.1)$$

where T is the kinetic and V the potential energy. Their expressions are derived from the properties of the physical system under study. The equations of motion are obtained from the Lagrangian using the Euler-Lagrange equations

$$\frac{d}{dt} \left(\frac{\partial L}{\partial \dot{q}} \right) = \frac{\partial L}{\partial q}, \quad (2.2)$$

where t is the time and q are the generalised coordinates of the system with their time derivatives \dot{q} . These yield the same equations of motion as Newton's laws, but their derivation is usually simpler and more succinct since the Lagrangian method does not require the enumeration of all the forces that are or can be at play in the physical system. It is not only for its efficiency or aesthetic value that we learn to appreciate this alternative approach, however: when those of us who long to get more familiar with the infinities of the universe, be they large or small, tackle quantum field theories, we find familiar mathematical constructs.

The Lagrangian density of the Standard Model of particle physics (SM) in one of its simplest forms, often reproduced on T-shirts, coffee mugs, and other basic necessities in a physicist's

life, reads as follows:

$$\begin{aligned}
 \mathcal{L} = & -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} \\
 & + i\bar{\psi}D\psi \\
 & + \psi_i y_{ij} \psi_j \phi + h.c. \\
 & + |D_\mu \phi|^2 - V(\phi).
 \end{aligned} \tag{2.3}$$

Compact as it is, it conceals considerable complexity within its notation. $F_{\mu\nu}$ stands for the vector field strength tensors of the photon, the gluon, and the weak bosons. The first term, therefore, describes the dynamics of those vector bosons and their interactions with each other (and with themselves). In the second term, ψ stands for the fermion fields, and D for their covariant derivatives (the slash means that they are contracted with the gamma matrices). Both the fermion dynamics and their interactions with the vector bosons are contained in these three symbols. Indeed, they encode such a diversity of processes that all those described later in this work are, at first order, derived from this second term. The third term describes the interaction of the fermions with the Higgs boson field, which gives them their masses. Finally, the fourth and fifth terms are again free of fermion fields: the fourth is the kinetic term of the Higgs boson, the only known fundamental scalar, and its interactions with the massive vector bosons. The fifth term is the Higgs potential, often likened to a Mexican hat, which allows for Higgs self-interactions.

The elementary particles of the SM are referred to as fermions and bosons grouping them by their spin, which is half-integer for fermions and integer for bosons. With this and only this distinction, there exist two elementary particles in the Standard Model. With the further distinction between vector (spin 1) and scalar (spin 0) bosons, there are three. The fermions come in two families: quarks and leptons. There are three generations of each family, and each generation has two components. For the quarks, these components are the up-type and down-type quarks. For the leptons, the components are the charged leptons and the neutrinos. Of the vector bosons, we distinguish the gluon, the W and Z bosons, and the photon. This brings the count of elementary particles up to seventeen: four vector bosons, one scalar boson, six quarks and six leptons. These seventeen can be arranged into more or less neat box configurations that conceptually stand in for the SM as a whole (another candidate design for T-shirts and coffee mugs) and the classification up until now has the advantage that each elementary particle has its own distinctive name, with the possible exception of the neutrinos. The flavour basis, in which the fermions have names, is not the only way to split the fermions into three groups (generations). The mass basis is a notable alternative, but changing the basis does not change the dimension, *i.e.* the number of groups. Still, we do not need to stop counting elementary particles here: most of the particles on the list of seventeen have a distinct antiparticle. Additionally, any particle with nonzero spin has two (for fermions and massless vector bosons) or three (for massive vector bosons) possible polarisation states. Finally, the quarks and gluons have colour charges: quarks can be red, blue,

2.1. The Standard Model of particle physics

or green (antiquarks can be anti-blue, anti-red, or anti-green) and gluons combine colour and anti-colour into nine charge states. One of those charge states is sterile, which means that it does not couple to any SM fields, and can therefore neither intervene in any detectable process, nor be detected by itself. Excluding sterile particles, the highest number of elementary particles we can arrive at for the SM is 118.

These particles enter bound states (a count of the number of composite SM particles will not be attempted here) scatter off each other, transform into one another, and interact in various other ways. The probabilities and dynamics of these interactions can be calculated using the Lagrangian density given in Eq. 2.3.

2.2 S-matrix elements in a hurry

Let us attempt a fast-tracked motivation of the calculation of the probability amplitudes we will soon need to use. No amount of mathematical rigorousness will even be claimed to be attempted in what follows, and sensitive readers, who do not need this motivation anyway, may wish to skip to Section 2.3.

Given an initial state i , we wish to calculate the probability of arriving at the final state f . Putting the f into a bra and the i into a ket, the probability is $P_{i \rightarrow f} = |\text{out}\langle f | S | i \rangle_{\text{in}}|^2$ (in bracket notation, it is customary to read from right to left). The relation of a set of final states $\{|f\rangle\}$ to a set of initial states $\{|i\rangle\}$ is described, in scattering theory, using the S -matrix. The operator which transforms an asymptotic “in-state” $|i\rangle_{\text{in}}$ at $t = -\infty$ to an asymptotic “out-state” $|i\rangle_{\text{out}} = S|i\rangle_{\text{in}}$ at $t = \infty$ is also known as S : the S -matrix is a representation of this operator in the basis of states mentioned above, i.e.

$$S_{fi} = \text{out}\langle f | S | i \rangle_{\text{in}}. \quad (2.4)$$

The interesting part of the S -matrix is the one where $f \neq i$, that is, where something changes or happens. To this end, S can be written as

$$S = \mathbb{1} + 2iT, \quad (2.5)$$

where the operator T describes the interaction and the factor $2i$ is added for convenience.

To calculate T , we need to bring the states $|i\rangle$ and $\langle f|$ from their far-infinite times to a moment where some interaction can take place. The time-evolution of the states is governed by the time-dependent Schrödinger equation with the operator \hat{H} for the Hamiltonian, which can be obtained from the Lagrangian with a Legendre transformation

$$H = \frac{\partial L}{\partial \dot{q}} \dot{q} - L,$$

so that the time-dependent Schrödinger equation is

$$i\hbar\partial_t |\Psi\rangle = \hat{H} |\Psi\rangle.$$

Since all mathematical rigorousness is abandoned, this differential equation is easily solved:

$$|\Psi(t_1)\rangle = e^{-i\hat{H}(t_1-t_0)/\hbar} |\Psi(t_0)\rangle.$$

We can split \hat{H} into a “free” part and an “interaction” part:

$$\hat{H} = \hat{H}_0 + \hat{H}_{\text{int}}.$$

The free part \hat{H}_0 describes a world where no interactions between fields take place, the number of particles always remains constant, and states do not change as they evolve. It turns out that this free world is the one where we can cleanly define the in- and out-states, so that they are not always interacting and changing. This is what is meant when they are called asymptotic.

The point is now to evolve the asymptotic “free” state $|i\rangle$ to current time, briefly switch on the interaction \hat{H}_{int} , and then evolve the resulting “interacted” state according to the free theory to an asymptotic out-state and see whether (or how much) it matches state f . The evolution of an asymptotic state, which is an eigenstate of the time-independent Schrödinger equation

$$\hat{H}_0|\Psi\rangle = E\Psi|\Psi\rangle,$$

is just a phase $e^{-i\hat{H}_0 t/\hbar}$, which will cancel by conservation of energy when the interacted state is evolved to an out-state again. Similarly to the time evolution, the interaction involves $e^{-i\hat{H}_{\text{int}} t/\hbar}$, but since we only switch it on for a short time Δt , we can expand the exponential to first order in Δt :

$$e^{-i\hat{H}_{\text{int}} t/\hbar} = \mathbb{1} - i\hat{H}_{\text{int}}\Delta t + \mathcal{O}(\Delta t^2).$$

Putting it all together, ${}_{\text{out}}\langle f|i\rangle_{\text{in}} = \langle f|S|i\rangle \approx \langle f|\mathbb{1} - i\hat{H}_{\text{int}}|\mathbb{1}\rangle$ and we can identify the transition matrix element T_{if} with the amplitude $\langle f|\hat{H}_{\text{int}}|\mathbb{1}\rangle$ where we have introduced natural units $\hbar = c = 1$ and where any coefficients dropped or added are swept under the rug of this not being a real derivation. The sensitive reader may now resume.

2.3 The photon polarisation in $B \rightarrow K\pi\pi\gamma$ decays

In order to simplify the notation, we will only refer to B^- mesons (which contain a b quark) in the discussion that follows, although the results will apply just as well, after applying charge-parity conjugation, to B^+ mesons (which contain a \bar{b} quark). B^- mesons are pseudoscalar particles, which means that their total spin J is 0 (the “pseudo” refers to the fact that, unlike true scalars, they have negative parity). In this work, we examine decays of these mesons into a kaonic resonance K_{res}^- and a photon. The K_{res}^- then decays to give three charged hadrons K^- , π^+ , and π^- in the final state. These decays are interesting because of their flavour-changing neutral current (FCNC) process of a bottom quark transitioning to a strange quark and a photon. In the SM, this transition can only happen through a loop such as the one shown




Figure 2.1 – Feynman diagram of a FCNC $b \rightarrow s\gamma$ transition through a weak loop [15].

in Figure 2.1. In the B^- rest frame, the K_{res}^- and γ have equal and opposite momenta, and the total orbital angular momentum \vec{L} of the system is the sum of the two daughters’ orbital angular momenta, so $\vec{L} = \vec{r}_\gamma \times \vec{p}_\gamma + \vec{r}_{K_{\text{res}}^-} \times \vec{p}_{K_{\text{res}}^-} = (\vec{r}_\gamma - \vec{r}_{K_{\text{res}}^-}) \times \vec{p}_\gamma$ is perpendicular to the photon momentum \vec{p}_γ (and the K_{res}^- momentum $\vec{p}_{K_{\text{res}}^-}$). Defining the z axis parallel to these momenta, the projection of \vec{L} on the z axis is zero. Since the B^- has spin 0, conservation of angular momentum requires that the spin projections of the K_{res}^- and the γ on the z axis be equal in magnitude and opposite in sign. The photon is a vector boson (spin 1), but because it is massless, it has no spin 0 projection. There are two allowed configurations: the photon spin can be aligned with its momentum (“right-handed”) or anti-aligned (“left-handed”). The K_{res}^- must consequently have spin projection ± 1 in the opposite direction.

The aim of the amplitude analysis discussed in this work is to measure the relative proportions of these two configurations in $B^- \rightarrow K^-\pi^+\pi^-\gamma$ decays. The SM predicts that the photon is mostly left-handed. This is due to an interesting property of the weak interaction, which in the SM mediates FCNC processes such as the $b \rightarrow s\gamma$ transition at the heart of the $B^- \rightarrow K_{\text{res}}^-\gamma$ decay: the weak interaction couples only to left-chiral particles and right-chiral antiparticles. For $b \rightarrow s\gamma$, by a similar argument as for $B^- \rightarrow K_{\text{res}}^-\gamma$, the spin 1 of the photon and the spin $\frac{1}{2}$ of the s quark must be in opposite directions in order to sum to the spin $\frac{1}{2}$ of the initial-state b . This implies opposite spins for the b and s quarks. The relationship between helicity (spin projection on momentum) and chirality (how a particle transforms between inertial reference frames) is mathematically complex; both attributes can be described using “handedness,” but one can only entirely identify right-handed chirality with right-handed helicity for massless particles. If a particle is massive, there always exists a reference frame where its helicity is

“opposite” its chirality. For the present purposes, let it suffice to say that flipping a quark’s chirality “costs” a factor equal to the quark’s mass, as it is the mass term that mixes the two chiralities. Since $\frac{m_s}{m_b} \approx 0.02$, right-handed photons are suppressed in the SM with respect to left-handed photons in $b \rightarrow s\gamma$ transitions.

It is worth pointing out again that this suppression is due to two factors in the SM: firstly, only the weak interaction can mediate FCNC transitions, and secondly, the weak interaction is maximally parity-violating, *i.e.* it couples exclusively to particles of a specific chirality. The other fundamental forces of the SM do not have this property. One can thus easily imagine a non-SM interaction which does not violate parity, or which violates parity in a way different from that of the weak interaction. A contribution from such an interaction to the $b \rightarrow s\gamma$ process could introduce a shift in the relative probabilities of the two photon polarisations with respect to the SM prediction. In probing for New Physics (NP), it is convenient to describe $b \rightarrow s\gamma$ without explicitly involving the SM interactions, and instead grouping possible contributions to the process based on their mathematical structure, which reflects their physical properties. The method chosen to achieve this is the operator product expansion (OPE) where the long-distance part of a process, like the initial and final states, is contained in the matrix elements of the effective operators, while the short-distance physics, like the interactions that mediate the transition, is described by coefficients, known as Wilson coefficients, that determine the relative importance of the operators.

The number of local operators that contribute to a $b \rightarrow s\gamma$ transition is limited. They include the current-current operators $\mathcal{O}_{1,2}$, the QCD penguin operators \mathcal{O}_{3-6} , the electromagnetic penguin operator \mathcal{O}_7 , and the chromomagnetic penguin operator \mathcal{O}_8 .

The effective Hamiltonian [16] is built from these operators and their Wilson coefficients:

$$\mathcal{H}_{\text{eff}} = -\frac{4G_F}{\sqrt{2}} V_{ts}^* V_{tb} \sum_{i=1}^8 \mathcal{C}_i(\mu) \mathcal{O}_i(\mu) \quad (2.6)$$

where G_F is the Fermi constant, V_{ts} and V_{tb} are elements of the Cabibbo-Kobayashi-Maskawa (CKM) matrix, which links the quarks’ mass and flavour eigenstates, and μ is the energy scale at which the OPE is performed. While μ does not show up in the expressions for the local operators, the scale chosen influences their form, *e.g.* through the inclusion or omission of some quark flavours.

The Wilson coefficients $\mathcal{C}_i(\mu)$ are calculated from perturbation theory. In order to remove a dependence of these calculations on the regularisation scheme, one defines effective coefficients $\mathcal{C}_7^{\text{eff}}(\mu)$ and $\mathcal{C}_8^{\text{eff}}(\mu)$ which are the sum of $\mathcal{C}_7(\mu)$ or $\mathcal{C}_8(\mu)$ respectively and a linear combination of the coefficients $\mathcal{C}_1(\mu), \dots, \mathcal{C}_6(\mu)$ [17].

The main contribution to $b \rightarrow s\gamma$ comes from the matrix elements of the electromagnetic penguin operator \mathcal{O}_7 , which include the interaction shown in Figure 2.1. The electromagnetic

penguin operator

$$\mathcal{O}_7 = \frac{e}{16\pi^2} m_b \bar{s} \sigma_{\mu\nu} b F^{\mu\nu} \quad (2.7)$$

couples the s and b quarks to the electromagnetic field tensor $F_{\mu\nu}$ through the dipole tensor $\sigma^{\mu\nu}$ which can be written as a commutator of gamma matrixes: $\sigma^{\mu\nu} = \frac{i}{2} [\gamma^\mu, \gamma^\nu]$. \mathcal{O}_7 can be split by inserting a factor 1 in the form of the sum of two operators

$$1 = \frac{1 - \gamma_5}{2} + \frac{1 + \gamma_5}{2}, \quad (2.8)$$

which project the b quark onto its left-chiral and right-chiral component, respectively. We recall that the b chirality is directly linked to the photon helicity, so this amounts to splitting \mathcal{O}_7 by the photon polarisation:

$$\mathcal{O}_7^{(\prime)} = \frac{e}{16\pi^2} m_b \bar{s} \sigma_{\mu\nu} \frac{1 \mp \gamma_5}{2} b F^{\mu\nu}, \quad (2.9)$$

where the unprimed or primed operator describes the transition to a left-handed (L) or a right-handed (R) photon, respectively. The factor m_b could be absorbed into the Wilson coefficients but is here retained to simplify the comparison with the SM expectation, and we can define the Wilson coefficient \mathcal{C}'_7 for just the operator \mathcal{O}'_7 , while \mathcal{C}_7 belongs to \mathcal{O}_7 . We obtain the transition amplitudes for the decays $B^- \rightarrow K_{\text{res}}^- \gamma_L$ and $B^- \rightarrow K_{\text{res}}^- \gamma_R$ from the effective Hamiltonian:

$$\begin{aligned} \mathcal{A}_L(B^- \rightarrow K_{\text{res}}^- \gamma_L) &= \langle K_{\text{res}}^- \gamma_L | \mathcal{H}_{\text{eff}} | B^- \rangle, \\ \mathcal{A}_R(B^- \rightarrow K_{\text{res}}^- \gamma_R) &= \langle K_{\text{res}}^- \gamma_R | \mathcal{H}_{\text{eff}} | B^- \rangle. \end{aligned} \quad (2.10)$$

The operators \mathcal{O}'_7 and \mathcal{O}_7 can be transformed into each other by a parity transformation and a rotation. It follows (see Ref. [18] for details) that the amplitudes from just these two operators, $\langle K_{\text{res}}^- \gamma_L | \mathcal{O}_7 | B^- \rangle$ and $\langle K_{\text{res}}^- \gamma_R | \mathcal{O}'_7 | B^- \rangle$, are proportional to each other with a proportionality factor $P_{\text{res}} (-1)^{J_{\text{res}}-1}$, where P_{res} and J_{res} are the parity and spin of the kaonic resonance K_{res} , respectively. Note that the “crossed” terms $\langle K_{\text{res}}^- \gamma_L | \mathcal{O}'_7 | B^- \rangle$ and $\langle K_{\text{res}}^- \gamma_R | \mathcal{O}_7 | B^- \rangle$ are 0. The proportionality factor is just a sign, and thus the ratio of the Wilson coefficients $\frac{\mathcal{C}'_7}{\mathcal{C}_7}$ is equal, up to a sign, to $\frac{m_s}{m_b}$ in the SM.

What about the other operators in \mathcal{H}_{eff} , though? We have so far neglected to make a proper distinction between the processes $b \rightarrow s\gamma$ and $B^- \rightarrow K_{\text{res}}^- \gamma$. Since quarks can only be observed in bound states, the latter decay is what we are actually dealing with. Naïvely, it differs from $b \rightarrow s\gamma$ only by the presence of a \bar{u} quark, called “spectator” quark because it does not participate in the crucial $b \rightarrow s\gamma$ transition. The presence of the spectator quark, or, alternatively, the fact that $B^- \rightarrow K_{\text{res}}^- \gamma$ is really a hadron decay and not a quark decay, rather complicates the picture: corrections from QCD, notorious for being the most difficult of the SM interactions to calculate predictions from, must be taken into account. Thus, the decay amplitudes in Eq. 2.10 contain the dominant terms from \mathcal{O}_7 and \mathcal{O}'_7 as well as hadronic

2.3. The photon polarisation in $B \rightarrow K\pi\pi\gamma$ decays

corrections h_L and h_R [19]:

$$\begin{aligned}\langle K_{\text{res}}^- \gamma_L | \mathcal{H}_{\text{eff}} | B^- \rangle &= -\frac{4G_F}{\sqrt{2}} V_{tb} V_{ts}^* \mathcal{C}_7^{\text{eff}} g_{\text{res}}(0) + h_L, \\ \langle K_{\text{res}}^- \gamma_R | \mathcal{H}_{\text{eff}} | B^- \rangle &= -\frac{4G_F}{\sqrt{2}} V_{tb} V_{ts}^* \mathcal{C}_7'^{\text{eff}} P_{\text{res}}(-1)^{J_{\text{res}}-1} g_{\text{res}}(0) + h_R,\end{aligned}\quad (2.11)$$

where $g_{\text{res}}(0)$ is a hadronic form factor for the $B^- \rightarrow K_{\text{res}}^- \gamma$ transition at $q^2 = 0$.

While we are on the subject of hadrons, let us not forget that K_{res}^- is not in the final state of our $B \rightarrow K\pi\pi\gamma$ decay: its decay into $K^-\pi^+\pi^-$ must also be described. That not enough, our generic “kaonic resonance” can be one of many resonances with different spin and decay properties. This abundance of decay modes will actually turn out to be crucial to the sensitivity of the $B \rightarrow K\pi\pi\gamma$ decay to the photon polarisation. The other crucial factor is the presence of three hadrons in the final state.

The differential decay rate for the full decay can be written as [18]

$$\begin{aligned}d\Gamma(B^- \rightarrow K_{\text{res}}^- (\rightarrow K^-\pi^+\pi^-) \gamma) = & \left| \sum_i \mathcal{A}_L \left(B^- \rightarrow K_{\text{res},i}^- \gamma_L \right) \mathcal{A}_L \left(K_{\text{res},i}^- \rightarrow K^-\pi^+\pi^- \right) \right|^2 + \\ & \left| \sum_i \mathcal{A}_R \left(B^- \rightarrow K_{\text{res},i}^- \gamma_R \right) \mathcal{A}_R \left(K_{\text{res},i}^- \rightarrow K^-\pi^+\pi^- \right) \right|^2,\end{aligned}\quad (2.12)$$

where the sums go over all allowed kaonic resonances $K_{\text{res},i}^-$ and the amplitudes from Eqs. 2.10 and 2.11 have acquired process-dependent indices in the expressions for parity, spin, hadronic form-factor, and hadronic correction. The terms for right-handed (R) and left-handed (L) photons are summed incoherently since they do not interfere due to the photon polarisation being an observable quantity. The hadronic form factors, while process-dependent, are equal for both helicities. If we neglect the small hadronic corrections h_L and h_R (see Ref. [20] for a discussion of the magnitude of these long-distance QCD corrections), we can pull the Wilson coefficients \mathcal{C}_7 and \mathcal{C}_7' out of the sums. That way, the expression for the differential decay rate can show a dependence on the photon polarisation parameter

$$\lambda_\gamma = \frac{|\mathcal{C}_7|^2 - |\mathcal{C}_7'|^2}{|\mathcal{C}_7|^2 + |\mathcal{C}_7'|^2}. \quad (2.13)$$

The hadronic decay amplitudes are described in the isobar formalism as consecutive two-body decays such as $K_{\text{res},i}^- \rightarrow K^- R_j (\rightarrow \pi^+\pi^-)$ or $K_{\text{res},i}^- \rightarrow R_j (\rightarrow K^-\pi^+) \pi^-$, introducing further intermediate resonances R_j of two final-state hadrons. Each such decay chain is treated as a separate “amplitude” and we can write

$$d\Gamma \propto \frac{1+\lambda_\gamma}{2} \left| \sum_k a_k e^{i\phi_k} A_{k,L} \right|^2 + \frac{1-\lambda_\gamma}{2} \left| \sum_k a_k e^{i\phi_k} A_{k,R} \right|^2. \quad (2.14)$$

The new index k runs over the decay chains. The relative magnitudes and phases of the amplitudes are given by a_k and ϕ_k respectively. This leaves the task of describing the “normalised” amplitudes $A_{k,L}$ and $A_{k,R}$. Dropping constant factors that are equal between the left-handed and right-handed terms, the weak decay amplitude reduces to a spin factor $S_{i,L/R}$ that absorbs the relative sign $P_{\text{res},i}(-1)^{J_{\text{res},i}-1}$ and a normalised Blatt-Weisskopf coefficient B_{L_B} [21] that depends on the relative angular momentum L_B and the breakup momentum q_B (the absolute value of the three-momentum of the $K_{\text{res},i}^-$ or the photon in the B^- rest frame). The strong decays are described by resonance propagators \mathcal{T}_i^k and \mathcal{T}_j^k for the three- and two-hadron resonances $K_{\text{res},i}^-$ and R_j in the decay chain, respectively, and spin factors $S_{i,j,L/R}^k$ that depend on the resonances’ spin structures:

$$A_{k,L/R}(\mathbf{x}) = B_{L_B}(q_B(\mathbf{x}), 0) S_{i,L/R}(\mathbf{x}) \mathcal{T}_i^k(\mathbf{x}) \mathcal{T}_j^k(\mathbf{x}) S_{i,j,L/R}^k(\mathbf{x}). \quad (2.15)$$

These expressions depend on the four-momenta of the final-state particles, \mathbf{x} . It is worth noting that the magnitudes a_k and phases ϕ_k are presumed to be equal for the left-handed and right-handed cases, as the strong decays of the hadronic resonances are invariant under parity transformations. The difference between $A_{k,L}$ and $A_{k,R}$ lies in the spin factors $S_{i,L/R}$ and $S_{i,j,L/R}^k(\mathbf{x})$. Interferences between decay channels with different resonances ensure that λ_γ does not cancel in the differential decay rate, and thus, that the $B \rightarrow K\pi\pi\gamma$ decay is sensitive to λ_γ .

Let us now return to the matter of charge-parity conjugation, that is, to B^+ decays. Following the same argument structure as above, we can show that the photon emitted in $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays should be predominantly right-handed in the SM, since now the quark transition in question is $\bar{b} \rightarrow \bar{s}\gamma$, and the weak interaction requires the \bar{b} and \bar{s} anti-quarks to be right-chiral. By changing the sign of the charges and three-momenta of all the particles involved, we can treat B^- decays like B^+ decays and thus perform a “ CP -averaged” measurement.

Chapter 3

The LHCb spectrometer

Experimental research in the natural sciences involves testing the validity of theoretical predictions in the laboratory. Experimental particle physics is no exception. The theoretical predictions we want to test usually come from the SM and concern the behaviour of the members of its particle zoo. To test the validity and applicability of the SM to as complete an extent as possible, we do not want to limit ourselves to the relatively small pool of stable particles. Unstable particles, however, are not so easily available for observation as stable ones. By Einstein's famous mass-energy equivalence ($E = mc^2$) [22], massive particles, including unstable ones, can be created from energy.

A widely accepted fundamental law of the universe is the conservation of (total) energy: energy can neither be created nor destroyed, cannot be spontaneously generated or vanish. Energy - including, by the above equivalence, matter - can only be transformed. Particles can liberate some or all of their energy for transformation by decaying or by interacting with other particles. Such a process has an initial state, usually one (in the case of a decay) or two (in the case of a collision) particles, and a final state of at least one (collision) or two (decay) particles. The process is characterised by specifying the four-momenta of all the initial and final state particles. Detectors are used to measure the four-momenta. Roughly speaking, high-energy processes are interesting for particle physicists because they allow the study of a larger number of unstable particles: a particle with mass m (rest energy mc^2) can only be created in a process where the available energy is at least mc^2 .

There exist natural sources of high-energy processes such as cosmic rays. In the early days of particle physics research, interactions of cosmic rays with particles in the Earth's atmosphere were the primary source of experimental knowledge about the behaviour of unstable particles. While cosmic rays are available "for free," there are several disadvantages to their experimental usage. The obvious disadvantages are related to their unpredictability: while we can know the energy spectrum of cosmic rays at the detector position quite well, it is impossible to predict when exactly the next high-energy process will occur, which kind of particle will be

Chapter 3. The LHCb spectrometer

the cause, and exactly how high its energy will be. The rate of high-energy processes is also generally so low that to record a sizeable sample of them would require a very large detection volume, a very long recording time, or both. Man-made sources of high-energy particles, such as (particle) accelerators, can yield much higher rates and are a much more dependable experimental source.

The highest-energy accelerator built so far is the Large Hadron Collider (LHC) at the European Organization for Nuclear Research, referred to as CERN (from the acronym of the French name of its predecessor council). The LHC is a 27 km ring collider that intersects the border between Switzerland and France near the Swiss city of Geneva. Together with its ancillary / preparatory accelerators (often, part of the structure of the most recent biggest collider on the site was used as the pre-accelerator and injector for the next big, or even higher-energy, collider project), the LHC accelerates bunches of protons and/or lead ions in two beams, clockwise and counter-clockwise, which are made to collide at several experimental points. When two bunches collide, there can be many proton-proton (or ion-proton, or ion-ion) collisions. In the LHC experiments, all collisions during a bunch crossing are referred to as one “event.” The period of operation of the LHC between 2009 and 2013 is known as Run 1. It was followed by a phase of maintenance, developments and upgrades before the second main period of operation, Run 2, which started in 2015 and concluded in 2018. During Run 1, proton-proton collisions at the LHC reached a centre-of-mass energy of 8 TeV. In Run 2, the centre-of-mass energy went up to 13 TeV, and in Run 3, which started in 2022, to 13.6 TeV.

The LHCb experiment sits at one of the LHC collision points in its experimental cavern at a depth of about 100 m (the whole LHC tunnel being underground) near the French town of Ferney-Voltaire just across the Swiss border. The researchers of the LHCb collaboration mostly study the physics of b and c quarks, and the detector has been designed for this purpose. It is a single-arm forward spectrometer covering a pseudorapidity range from $\eta = 2$ up to $\eta = 5$. “Forward” here means that it sees those particles created in the pp -collisions that travel in the forward direction, close to the proton beam line, as opposed to those that have most of their momentum perpendicular to the beam line. There is technically no “backward” direction, the colliding beams being symmetric, but LHCb only covers one of the two forward regions - otherwise it would be a double-arm forward spectrometer. The word “spectrometer” refers to the detector’s task of measuring particle masses and momenta and their distributions (spectra). This is achieved using multiple sub-detectors with different scopes of duties. The description of the detector that follows refers to the configuration in Run 1 and Run 2, during which the data used in the amplitude analysis were recorded.

A schematic view of the LHCb experiment is shown in Figure 3.1. The LHCb coordinate system is defined such that the origin is at the interaction point (IP), where the protons or ions from the two LHC beams are brought to collision. The z axis points from the IP into the forward region covered by the LHCb spectrometer (towards the right in Figure 3.1). The y axis points vertically upwards, and the horizontal x axis completes the right-handed coordinate system. The direction of motion of particles in the LHCb acceptance towards higher values of z is




Figure 3.1 – The LHCb detector. Particles created at the IP in the VELO on the left-hand side pass downstream (towards the right) through the systems of sub-detectors.

called “downstream”, and the opposite direction is “upstream”. The two regions of the LHCb cavern on either side of the spectrometer are known as A-side (for “access”, $x > 0$) and C-side (for “cryo”, $x < 0$). Figure 3.1 is a view of the spectrometer from the C-side.

The process of detecting and reconstructing the path taken by a charged particle is known as tracking. Charged particles leave signals as they pass through layers of tracking detectors (trackers) with air or vacuum between them. The tracking layers are perpendicular to z and segmented so that, for each signal, the transverse coordinates x and y can be reconstructed at the z position of the layer. The tracking sub-detector that is closest to the IP is the Vertex Locator (VELO). It consists of two arrays of roughly semicircular silicon-strip modules perpendicular to the beam, with a clearance of radius 8 mm at the centre of the circle through which the beam passes during operation. The outer radius is 42 mm, so a VELO module looks a bit like a Mini-CD that has been broken in half. A photograph of one of the VELO arrays taken during assembly is shown in Figure 3.2. Each module actually covers a little more than half of the azimuthal angle, so the two halves on each side of the beam can overlap and thus avoid a dead zone. There are two kinds of VELO modules: one is called R-sensor because it measures the distance of a signal to the centre with circular strips, and the other ϕ -sensor, which measures the azimuthal angle with radial strips. To keep the LHC beam vacuum intact, the VELO arrays, while mounted inside the vacuum pipe, are separated from the beam vacuum by corrugated sheets of aluminium, 0.3 mm thick, that also protect the VELO from radio-frequency beam noise. The two sides of the VELO can be retracted from the beam,



Figure 3.2 – One half of the VELO sub-detector, with its semicircular silicon modules.

which needs a higher clearance during ramp-up, and closed again when the beams are stable. Bringing the VELO sensors as close to the beam as possible allows high impact-parameter resolution of tracks near the collision point.

The next tracking station is the Tracker Turicensis (TT). It is a single planar tracking station that is 150 cm wide and 130 cm high, placed perpendicular to the beam line between the VELO and the magnet. Downstream of the magnet are another three tracking stations, each of which has an Inner Tracker (IT) and an Outer Tracker (OT) component. The inner region, close to the beam pipe, requires better spatial resolution of track hits than the outer region to achieve comparable track reconstruction efficiency and relative momentum resolution between the two regions. This is due to the fact that the trajectories of high-momentum particles are less affected by the magnetic field, and the fact that the number of charged particles per unit area is highest close to the beam pipe. There are significantly more tracks per collision, on average, that pass through the inner region, and these tracks tend to come from more energetic particles than those in the outer region. The IT and the TT, the latter of which is much closer to the IP where the transverse area of the LHCb acceptance is much smaller than for the downstream tracking stations, use silicon microstrip sensors. The sensors are arranged, for each of the four silicon tracker stations (one TT station and three IT stations), in four layers that are rotated with respect to one another. The first and fourth layers have vertical strips, while the two layers on the inside are rotated by $+5^\circ$ and -5° from the vertical around the z axis. This configuration produces a single-hit resolution just above 50 μm with silicon strip pitches of about 200 μm . The OT only needs to provide a spatial resolution of approximately 200 μm , which is achieved using straw-tube detectors with an inner radius of 4.9 mm. The straw tubes are arranged in four modules per station. Each module has two layers of tubes, and the four modules are arranged in the same pattern of vertical – off-vertical – off-vertical – vertical orientations as the silicon tracker stations.

The LHCb magnet is a warm dipole magnet positioned between the TT and the downstream tracking stations [23]. Two coils sit above and below the beam line, surrounded by a 1450 ton




Figure 3.3 – The LHCb dipole magnet, showing the constituent iron plates of the yoke and bent pancakes of the coils. Current connections to the coils and water cooling pipes are shown, as well as some of the yoke support structure. Distances are indicated in mm [24].

window-frame iron yoke with a wedged aperture. The predominantly vertical field bends the trajectories of charged particles in the horizontal plane. A schematic drawing of the magnet yoke and coils from the design stage of the magnet system is shown in Figure 3.3. Each of the two coils is made from 15 stacked layers of trapezoidal race-track shaped “pancakes.” Each pancake is a spiral coil made from an extruded aluminium Al-99.7 conductor with a length of about 300 m and a square 50×50 mm 2 cross-section. The conductor has a central duct of 24 mm diameter for water cooling. The tapering sides of the pancakes are bent by 45° along lines with a 250 mrad angle with the symmetry axis of the trapezoid. Five triplets of pancakes make up one coil [25]. The yoke is constructed of low-carbon steel (EN S235JR) plates of 80 to 100 mm thickness. The horizontal (vertical) aperture of the yoke window is parallel to the 300 mrad (250 mrad) acceptance boundaries, but with an additional 100 mm (600 mm) clearance. Iron “shims” are added on the sides of the pole faces to increase the field uniformity [26]. Beyond the shims, the yoke has cut-outs where the coils are mounted. The nominal current in the conductors of 5850 A has to be multiplied by 225 to obtain the total nominal current in each coil [27]. The field produced by the magnet is mostly vertical, with a maximal field strength of 1.1 T and a bending power of 4 Tm. The field polarity is periodically reversed by inverting the current direction to help cancel out detection asymmetries of charged particles. The polarity where the y component of the field is positive is referred to as “Up” or “MagUp” polarity, and the configuration where the y component is negative has the field pointing down, and is therefore referred to as “Down” or “MagDown.” The bending power is calculated by integrating the field along a path, in this case along the beam axis where the symmetry of the magnet leads to the highest and most uniform vertical field. It gives a measure

of the deviation of a charged particle's trajectory as it passes through the field. This deviation is used to measure the particle's momentum in the bending plane. The magnetic field in the LHCb acceptance must be precisely known to achieve accurate momentum reconstruction. The mapping of the LHCb dipole field for Run 3 of the LHC is described in Chapter 4.

As a charged particle passes through a dielectric medium, its electric field polarises the medium around it. If the speed v_p of the charged particle exceeds the speed of light in the medium $v_n = \frac{c}{n}$, where c is the speed of light in vacuum and $n \geq 1$ the refractive index of the medium, the wave fronts caused by the depolarisation of the regions the particle has just left interfere constructively to a conical wavefront that follows the particle, radiating light at an angle θ_c to the particle direction, with $\cos \theta_c = \frac{v_n}{v_p}$. This light emission is known as the Cherenkov effect. It can be used to measure the velocity of highly energetic charged particles. Together with measurements of the momentum and/or the energy of the particle, the mass can be calculated, which amounts to identifying the particle type. The LHCb detector has two sub-detectors that make use of Cherenkov light to provide particle identification (PID) information. For a charged particle that passes through a short dielectric medium, the conical geometry of the emitted light leads to an image of a ring on a screen or detector placed perpendicularly to the particle trajectory. The two sub-detectors are called RICH1 and RICH2, where RICH stands for Ring Imaging Cherenkov detectors. The RICH1 is located between the VELO and the TT, and covers low-momentum ($1 - 60 \text{ GeV}/c$) charged particles in the full LHCb acceptance, while the RICH2 sits between the downstream tracking stations and the calorimeters and covers high-momentum ($15 - 100 \text{ GeV}/c$) charged particles in the central (forward) region of the acceptance. Both sub-detectors use gas as the dielectric medium (radiator): RICH1 uses C_4F_{10} , which has a refractive index of $n = 1.0014$ at a photon wavelength λ of 400 nm, and RICH2 uses CF_4 , which has $n = 1.0005$ at $\lambda = 400 \text{ nm}$. The Cherenkov photons are reflected out of the LHCb acceptance by two sets of mirrors, one spherical and one flat, to segmented photon detectors. The radius of the reconstructed ring gives the opening angle θ_c of the Cherenkov cone. The PID performance of RICH1 is illustrated in Figure 3.4, where different types of particles can be identified by their highly separable distributions in $p - \theta_c$ space.

The calorimeter system measures the energy of charged as well as neutral particles. There are two basic systems: the electromagnetic calorimeter (ECAL), which measures the energies of photons and electrons and, downstream of it, the hadronic calorimeter (HCAL), which measures charged and neutral hadrons.

Before a particle enters the main ECAL structure, it passes through two sub-detectors that provide PID information: the Scintillator Pad Detector (SPD) and the Preshower (PS) detector. The SPD/PS system consists of two scintillation layers with a 15 mm lead converter plate sandwiched between them. Electrons, charged hadrons, and photons have different signatures in terms of the signals they leave in the SPD and PS. The SPD acts almost as a tracking station, in the sense that charged particles leave only a minimal amount of energy as a sign of their passage through the scintillator. Charged hadrons do not interact much in the thin lead converter plate so that they pass through the PS (and the ECAL) with no more than those




Figure 3.4 – Reconstructed momentum and Cherenkov angle in RICH1 for isolated tracks [28].

minimal energy deposit signals. Electrons, however, are induced by the lead plate to produce a cascade of secondary particles, each with a fraction of the electron's energy. The signal of such an electromagnetic (EM) shower in the PS is then much higher than the minimal signal left by a charged hadron. Photons have no electric charge and therefore leave no signal in the SPD, but like electrons, they shower in the lead plate. The EM showers of electrons and photons proceed into the ECAL proper, which is a “shashlik” structure of alternating lead absorber plates and scintillator tiles. Sixty-six such lead-scintillator layers make up one 42 cm long ECAL module. The length of the modules was chosen so that all EM showers are entirely contained within the ECAL, for optimal energy resolution. The transverse size of the square module faces is 40.4 mm in the inner region of the ECAL close to the beam line, 60.6 mm in the middle region, and 121.2 mm in the outer region. The SPD and PS scintillator pad sizes match those of the ECAL. The lateral segmentation of the sub-detectors is finer towards the beam line due to the much higher hit frequency per area in the central region. Neutral pions decay into pairs of photons which induce EM showers in the ECAL. Besides localising the energy deposits in the transverse plane, the lateral segmentation of the sub-detectors allows for some discrimination power between single photons and pairs of photons from neutral pions through comparison of the transverse shower distributions, see also Section 5.1.

Like the ECAL, the HCAL consists of layers of absorbers and scintillators, but unlike the ECAL, these layers are arranged parallel to and not perpendicular to the beam axis. The HCAL uses iron as the absorber material. Its transverse cell size is 131.3 mm in the inner region and 262.6 mm in the outer region. Hadrons are less likely to produce cascading particle showers than electrons and photons, which is why the HCAL is positioned after the ECAL. A hadronic shower has an EM component and a hadronic component, the latter coming from strong interactions with the shower medium. The variable relative energy loss through EM and strong processes makes it difficult to resolve with low uncertainty the energy of a hadronic shower. The required HCAL energy resolution is only moderate compared to the required resolution in




Figure 3.5 – The LHCb calorimeter system [29]. One half of the HCAL, PS and SPD are shown in open position, away from the beam axis.

the ECAL, however. Nevertheless, with a depth of 1.2 m, the HCAL is much longer than the ECAL.

In both ECAL and HCAL, the light produced by charged particles in the scintillators is read out using wavelength-shifting fibres and photodetectors. The light yield is proportional to the energy deposited in the scintillator. Figure 3.5 is a schematic view of the calorimeter sub-detectors, showing the different granularities of the lateral segmentation.

The electron's heavier cousin, the muon, does not induce any shower in the ECAL, and indeed leaves only the minimal energy deposits of a charged particle due to ionisation in the calorimeter sub-detectors (as well as in the tracking stations). This qualitative difference between electrons and muons comes from the fact that the muon mass is ~ 200 times that of the electron. An electron shower starts with the emission of a high-energy photon due to Bremsstrahlung in the electric field of a nucleus. The energy loss from Bremsstrahlung is inversely proportional to the squared mass, which explains why muons can pass through the LHCb spectrometer without significant energy loss. This is a boon when it comes to the identification of this particle type: at the downstream end of the spectrometer sit the muon stations, which are multi-wire proportional chambers, interspersed with 80 cm thick iron absorber plates. One extra muon station without iron absorber sits in between the downstream tracker and the ECAL, where it improves the transverse momentum resolution of the muon

system. In its central region, the radiation rate is so high that it would degrade the life-time of a wire chamber system to an unacceptable extent, so triple-GEM (gas electron multiplier) detectors are used there instead. Muons with a momentum above $3 \text{ GeV}/c$ traverse at least the first three muon stations, while those above $6 \text{ GeV}/c$ traverse all five. The PID efficiency for muons is around 97% for a 1 – 3% probability that charged pions are mistaken for muons.

The LHC delivered pp -collisions at a rate of 40 MHz. It is impractical to store all the signals collected in the many sub-systems of the LHCb detector at this high rate: the disk size of one such event is about 50 kB, and both bandwidth and storage limitations preclude this sort of greedy data collection. Only a small fraction of collisions lead to processes that are of interest to the LHCb physics programme anyway, such as $b\bar{b}$ pair production. And in only 15% of such cases are all the decay products of at least one of the beauty hadrons contained within the LHCb acceptance [30]. Interesting events should be stored for analysis, while the rest can be discarded: this is the task of the LHCb trigger system. It consists of three stages of increasing complexity in the selection requirements and decreasing pass rates. The L0 (level 0) trigger is hardware-based and selects events with high p_T/E_T signatures from the calorimeter system and the muon stations. Two subsequent high-level trigger stages, HLT1 and HLT2, perform (partial) event reconstruction. The trigger architecture was upgraded for Run 2, and notably, full event reconstruction was implemented in HLT2. The trigger system yielded event rates of 5 kHz in Run 1 [31] and 12.5 kHz in Run 2 [32].

Chapter 4

The LHCb magnetic field map

A precise knowledge of the magnetic field generated by the LHCb dipole is crucial for accurate reconstruction of the momenta of charged particles, which are deflected by the Lorentz force. The magnetic field strength in the detector region and its environment is also important from the perspectives of health and safety as well as detector operation: magnetic fields can disturb currents in electronic components of the detector. Specifically, the operation of photomultiplier tubes in the RICH sub-detectors requires that they be shielded from fields with strengths above the mT level [33]. There are, broadly, two ways to obtain quantitative information about the magnetic field. One uses simulations of the magnet and its environment. The other consists in measuring the field using Hall probes. These two sources of information are discussed in the following.

Hall probes are sensors that measure the magnetic field strength perpendicular to a thin strip of conductor thanks to the Hall effect. A magnetic field causes a deviation of the current in the conductor from a straight path to one side of the strip, resulting in a voltage drop between the two sides of the strip that is proportional to the magnetic field strength perpendicular to the conductor. Three such sensors arranged mutually perpendicular to each other allow the measurement of three perpendicular components of the magnetic field. In the following, the term “Hall probe” refers to such a configuration of three Hall sensors. For a map of the field in space, both the orientation and the location of the Hall probe with respect to some reference coordinate system must be known, which generally requires careful positioning as well as monitoring of the probe and its support apparatus. The three field components are usually interpreted as a measurement of the three-vector \vec{B} at some point in space \vec{x} , but since the sensors have nonzero width and cannot intersect, this space-vector \vec{x} is an approximation; in reality, the three sensors, while placed close together, measure their field components at slightly different points in space, and the measurements are technically averages over the surface of the conductor strip. The accuracy of the measurement of $\vec{B}(\vec{x})$ also depends on the accuracy of the measured location and orientation of the probe, and on the extent to which the three sensors’ orientations overlap (in the ideal case of perfectly perpendicular sensors, there

is no overlap). Another technical difficulty in the field measurements comes from the fact that the LHCb field map must cover a large volume (the region in the acceptance between $z = 0$ and $z = 12$ m alone has a volume of over 45 m³) with high granularity, which would require field measurements at thousands of points in space, some of which are hard to access and do not permit the installation of a measurement apparatus. It is therefore impractical to build a field map from measurements alone.

The magnetic field generated by a current in a coil and shaped by the surrounding materials can be calculated from Maxwell's equations. The equations' solutions determine the field at every point in space, so this approach is not hampered by measurement logistics. For a set-up as complex as the LHCb dipole and its environment, a discretised solution of the systems of partial differential equations is preferable to the intractability of an analytic approach. Discretisation, however, means that the resulting magnetic field values are necessarily approximations of the exact analytic solution, and while the discrepancy between the two can be reduced by decreasing the size of the discretisation steps, doing so will exhaust the available computing power before the step size is infinitesimal. The other limitation of the simulation approach lies in the correspondence between the simulation model and its physical counterpart, in this case the LHCb dipole and its environment. The model should be comprehensive and the positions, sizes, orientations, and magnetic properties of all its parts should match reality as precisely as possible, which is a manifestation of a rather poetic paradox: the most, the only accurate map of a territory is the territory itself, but as a map, it is useless [34]. In the present case, however, this cartographic problem illustrates the main virtue of field measurements, which is that they represent the magnetic field as it is in situ, without being affected by any shortcomings of modelling. Field values from both measurements and simulation are therefore used in concert to produce a map of the magnetic field of the LHCb dipole.

4.1 Overview of LHCb magnetic field maps

In the LHCb software, the field map is defined on a $10 \times 10 \times 10 \text{ cm}^3$ cubic grid that spans the volume where $-4 \text{ m} \leq x \leq 4 \text{ m}$, $-4 \text{ m} \leq y \leq 4 \text{ m}$, and $0.5 \text{ m} \leq z \leq 14 \text{ m}$. The map is split into quadrants by the $x = 0$ and $y = 0$ planes of the LHCb coordinate system. In the following, the term “map” shall refer to a full set of maps for all four quadrants.

The first map, produced in 2010 [35], was based on field measurements taken in 2005 [27] and a simulation of one quadrant of a model of the LHCb detector (the approximate mirror symmetry of the LHCb experiment with respect to the horizontal and vertical planes along the z axis was presumed exact), which included a description of the magnet coils and yoke as well as the iron absorbers in the HCAL and the muon stations, and the magnetic shieldings of the RICH1 and RICH2 sub-detectors. The field differences (residuals) between measured and simulated field values were parametrised using polynomials of order up to four in 11 overlapping regions along the z axis and in all four quadrants simultaneously. These parametrised corrections were then added to the values obtained from the field simulations at the grid points. The simulated field values at the measurement positions were extrapolated from the grid using the ROOT fit function TMuDimFit [36] with polynomials of order 2. The residual parametrisations of the three field components are independent of each other. Two sets of residual parametrisations were obtained for the two magnet polarisations.

Subsequent field maps were produced in 2011 [37] and 2014 [38] and based on the parametrised 2010 map and on field measurements taken inside the LHCb acceptance in 2011 and 2014, respectively. In both map updates, the previous field map was shifted, rotated, and scaled to minimise the discrepancies between the new field measurements and the map. The shifts were of the order of 1 cm and the rotations were of the order of a few mrad. The scale factors differed from 1 at the sub-permill level.

For the update of the magnetic field map for Run 3 of the LHC, it was decided to rebuild the map from scratch by rerunning the simulation after adjusting the simulation model. Several methods of incorporating the results of a new field measurement campaign, which was undertaken in January 2021, were tested. The measurement data and comparisons to the simulated field values are discussed in Section 4.2. The simulation model and the adjustments made to it are described in Section 4.3. The interpolation of magnetic field values from a discrete three-dimensional grid to points anywhere in the volume used throughout these studies is described in Section 4.4. The field map chosen for Run 3 and its validation on detector data are described in Section 4.5.

4.2 January 2021 magnetic field map measurements

Field measurements were recorded in the LHCb cavern between January 27th and January 29th, 2021. To record the field values, 37 Hall probes were mounted on a support bar with nominal distances of 57 mm between two probes. At two of these nominal probe slots were placed reflective targets for survey measurements of the bar's position in space. Survey targets were also placed on two smaller bars installed perpendicularly to the bar, close to its ends. The bar was installed parallel to the x axis on a vertical adjustment slide, which itself was mounted on a rail to allow shifting the position of the bar in the z direction. Figure 4.1 shows a photograph of the bar in the lowest y position on the vertical adjustment slide as well as the rail that extends into the magnet volume. The magnet coils are visible in the foreground on the top right and bottom right as well as in the background. Figure 4.2 is a close-up photograph of one end of the probe bar, showing some of the Hall probes and survey targets. With this set-up, the bar was moved to different positions in y and z . At each position, the magnetic field values measured by the 37 probes, which were spaced along the x direction, were recorded. The positions of the six survey targets on the bar were also recorded for each measurement. Together with a detailed survey of the bar and the placement of the Hall probes on the bar [39], the survey measurements allowed a precise reconstruction of the position of each probe for



Figure 4.1 – The probe bar in the LHCb dipole magnet on January 28th, 2021, viewed from the C side of the RICH1 sub-detector (upstream of the magnet).



Figure 4.2 – Close-up of one end (A side) of the probe bar in the LHCb dipole magnet on January 28th, 2021. Four Hall probes are fully in the picture, and three survey targets are visible: two on the vertical bar and one in the third slot in the bar, between the Hall probes.

each measurement. The rotation of the bar with respect to the LHCb coordinate system was also calculated from the survey data for each measurement.

Data was taken at 29 different positions of the vertical adjustment slide in z and five (three) different vertical positions of the bar for magnet polarity “Down” (“Up”). The positions of the Hall probes at which field measurements were recorded span the region from $x = -1.08$ m to $x = 1.08$ m, $y = -0.52$ m to $y = 0.64$ m, and $z = 2.51$ m to $z = 6.02$ m.

Two successive corrections were applied to the field values recorded by the Hall probes. First, the three-vector \vec{B} was rotated to correct for the difference between the orientation of the bar and the LHCb coordinate system determined from the survey measurements. Then, because it was found that the field measurements were less smooth along x (along the bar) than expected, a data-driven method of correcting for rotations of the Hall probes relative to each other was developed, which is described in the following.

Consider a probe which is rotated with respect to the LHCb coordinate system by a small angle ϑ around the axis parallel to the z axis that passes through the centre of the probe. The dominant component of the magnetic field is along the y direction (this holds for all the points in which measurements were recorded). The probe which is oriented off-axis picks up some of this dominant component in its nominal sensor for B_x :

$$B_x^{\text{meas}} = \cos \vartheta B_x^{\text{true}} + \sin \vartheta B_y^{\text{true}}, \quad (4.1)$$

$$B_y^{\text{meas}} = \cos \vartheta B_y^{\text{true}} - \sin \vartheta B_x^{\text{true}}, \quad (4.2)$$

$$B_z^{\text{meas}} = B_z^{\text{true}}. \quad (4.3)$$

The deviations from the accurate LHCb coordinate system can be presumed to be small, so we can approximate the field strength measured by the x sensor as the true x component of the field with a small contribution from the true y component of the field. The correction is linear in B_y^{true} with a correction coefficient c_x : $B_x^{\text{meas}} = B_x^{\text{true}} + c_x B_y^{\text{true}}$. A similar consideration for a probe that is slightly rotated around an axis parallel to x through its centre gives $B_z^{\text{meas}} = B_z^{\text{true}} + c_z B_y^{\text{true}}$ with a correction coefficient c_z . We will not repeat this exercise for probes that are rotated around the y axis because the fact that $|B_y| \gg |B_x|, |B_z|$ in the whole measurement volume means that only rotations of the probes around the x and z axes can be corrected with this method. We will return to discuss this later.

For a row of probes along the x axis, each with its own rotation with respect to the nominal orientation, the deviation of the measured value of B_x (B_z) from the true value is, to first order, proportional to B_y . The approximate symmetry of the magnet with respect to the $y = 0$ plane means that horizontal field components from the upper and lower coils should cancel out, and therefore that B_x and B_z should be close to 0 in the region surrounding the symmetry plane. The field measurements for the y position that lies closest to the plane are therefore chosen for the determination of the rotation corrections. The true value of B_x (B_z) cannot be independently determined, but it should be smooth along x . We therefore choose to use third-order polynomial fits to the B_x and B_z measurement values as a function of x as proxies for the true values.

The two sub-figures on top of Figure 4.3 show the fitted polynomials for B_x and B_z for one of the measurement positions of the bar as black dashed lines, while the measurements they are fitted from are shown as blue circles. After obtaining the polynomial truth proxies, the residuals ($B_{x(z)}^{\text{meas}} - B_{x(z)}^{\text{fitted}}$) are plotted for each probe against the measured value of B_y . The linear dependence is clearly visible for all probes, see the left-hand sub-figure of Figure 4.4 as an example. For some of them, the correlation between the residuals and the measured B_y is near zero, see *e.g.* the right-hand sub-figure of Figure 4.4, indicative of a linear dependence with a slope close to 0. As the correction coefficient calculated for such a probe is then also very small, these probes with negligible rotation are not excluded from the correction procedure.

The first order corrections applied for each probe based on the correction coefficients c_x and c_z lead to

$$B_x^{\text{corr}} = B_x^{\text{meas}} - c_x B_y^{\text{meas}}, \quad (4.4)$$

$$B_y^{\text{corr}} = B_y^{\text{meas}} + c_x B_x^{\text{meas}} + c_z B_z^{\text{meas}}, \quad (4.5)$$

$$B_z^{\text{corr}} = B_z^{\text{meas}} - c_z B_y^{\text{meas}}. \quad (4.6)$$

These corrections leave the field magnitude $|\vec{B}|$ invariant up to first order in the correction coefficients. The largest absolute values for the correction coefficients found are on the order of 10^{-3} , so the higher-order changes in field magnitude of order 10^{-6} are negligible.

4.2. January 2021 magnetic field map measurements




Figure 4.3 – Measured field values with corrections for the bar rotations from the survey (blue circles) and field values after the corrections for relative rotations (orange diamonds) as a function of x for one of the measurements near the beam line ($y = 0.01$ m, $z = 3.97$ m) at polarity “MagDown”. The B_x and B_z plots on the top left and top right respectively also show the polynomial fits used to calculate the offsets (see for example Figure 4.4) as dashed black lines. On the bottom are B_y (left) and $|\vec{B}|$ (right, dominated by the B_y component). The magnetic fields are in units of Tesla.




Figure 4.4 – Residuals $B_x^{\text{meas}} - B_x^{\text{fitted}}$ (in Tesla) from the 3rd degree polynomial fit of B_x measurements for the probe at $x = 34$ cm (left) and the probe at $x = -80$ cm (right) as a function of the B_y values measured by the same probe. The result of the linear fit is overlaid and the slopes, which correspond to the correction coefficients c_x for the respective probes, are given. Since the correlation in the left plot is relatively large compared that in the right plot, the linear dependence is more apparent on the left. Note the different scales (by almost two orders of magnitude) on the vertical axes.

Chapter 4. The LHCb magnetic field map

The corrections, calculated on the basis of data taken near $y = 0$ with “MagDown” polarity, also lead to smoother field values along x in independent regions, as illustrated in Figure 4.5, which shows data recorded with “MagUp” polarity at the top of the measurement volume.

The fact that the mean of the correction coefficients c_x and c_z (8×10^{-8} and -1.6×10^{-7} respectively) over all 37 probes is close to 0 points to the main blind spot of this correction method: its basic assumption that the true values of the fields B_x and B_z can be approximated by smoothing over the ensemble of probes. With this approach, the method can only correct for rotations of the probes relative to their mean orientation. If this mean orientation itself is not aligned with the LHCb coordinate system, a global rotation affects the array of probes. This global rotation cannot be corrected for, or even detected, with the data-driven correction method detailed above. Such a global rotation could result from a defect in the probe manufacturing process or the mounting of the probes on the bar. It could only be detected through calibration of the bar in a well-known reference field. The other shortcoming of this correction procedure has already been mentioned: it cannot account for relative, let alone global, rotations around the y axis. Such a correction could be accomplished in LHCb’s own magnetic field by measurements taken with a vertical bar; however, the bar supports are not set up for such an arrangement and the vertical clearance of the coils at the upstream end of the magnet is barely above 2 m, which is close to the length of the bar. The corrected field measurements for “MagDown” are shown in Figures 4.6-4.8.




Figure 4.5 – Measured B_x , B_y , B_z , and $|\vec{B}|$ values (in Tesla) with corrections for the bar rotations from the survey (blue circles) and after the corrections for relative rotations (orange diamonds) as a function of x for $y = 0.64$ m and $z = 4.92$ m at polarity “MagUp”. The data shown here were not used in the determination of the correction coefficients.

4.2. January 2021 magnetic field map measurements




Figure 4.6 – B_x measurements taken in January 2021 with polarity “MagDown,” with rotational corrections applied for the orientation of the bar and for the rotations of the Hall probes with respect to each other.




Figure 4.7 – B_y measurements taken in January 2021 with polarity “MagDown,” with rotational corrections applied for the orientation of the bar and for the rotations of the Hall probes with respect to each other.

Chapter 4. The LHCb magnetic field map




Figure 4.8 – B_z measurements taken in January 2021 with polarity “MagDown,” with rotational corrections applied for the orientation of the bar and for the rotations of the Hall probes with respect to each other.

4.3 Simulations of the magnetic field

Simulated field values are obtained using OPERA [40], a multiphysics software suite that uses finite element discretisation for electromagnetic analyses. The field equations are solved in this finite element method (FEM) with a scalar potential formulation algorithm, which was formerly known as TOSCA [41]. A basic model of the detector was implemented in OPERA and used in the course of the design, commissioning, and initial performance tests of the LHCb experiment. This model includes the magnet coils and yoke, as well as the shielding structures of both RICH sub-detectors and the iron absorbers in the HCAL and the muon stations (see Section 3). In order to obtain more accurate field values, it has been improved in the course of this work to correspond more closely to the experimental environment of LHCb. Figure 4.9 shows a view of the updated simulation model in OPERA. The following adjustments were made:

- The simulation volume was quadrupled to cover all four quadrants of the detector, discarding the assumption of exact symmetry with respect to the horizontal and vertical planes through the z axis.




Figure 4.9 – View of the simulation model for the calculation of magnetic field values in the OPERA software suite. The model comprises the magnet yoke (wine ■), the coils and the current lines (pale grey □), the clamps that support the yoke (sand ■), the yoke support chariot (green ■), the RICH1 (cyan ■), the RICH2 (teal ■), the HCAL (indigo ■), the muon stations (purple ■), the bunker (rose ■), and the scaffolding (olive ■).

- The maximal mesh size in the central magnet region, where both the field and its variations are highest, was decreased from 20 cm to 5 cm.
- The current in the conductor blocks that make up the coils was set to the nominal value of 225×5850 A, where previously the current density had been equal in all conductor blocks. Due to the varying cross-sections of differently shaped blocks, this constant current density meant that the total current varied between the conductor blocks. A constant current for all blocks is a more accurate description of the magnet coils.
- The surfaces of the pole faces, *i.e.* the planes on the top and bottom of the yoke window, which shape the direction of the magnetic field, were adjusted based on survey measurements of the pole faces *in situ*.
- The coil geometry was adjusted to account for the slight “opening” of the coils when current is flowing with respect to their “magnet off” state, also based on survey results.
- A simplified description of the connectors that provide current to the coils was added.
- Several ferromagnetic structures in the magnet environment were added, in simplified form, to the model:
 - the structure supporting the yoke, referred to as the yoke support chariot;
 - the scaffolding platforms and stairs that surround the magnet on the upstream side of the cavern;
 - the clamps added in the period between Run 1 and Run 2 to provide additional support to the coils on the upstream and downstream sides of the yoke;
 - the steel reinforcements of the concrete structure known as the “bunker” on the downstream side of the magnet, underneath the LHCb acceptance region.

Notably, other reinforced concrete structures in the LHCb cavern, such as the VELO alcove and the cavern walls and floor themselves were not included in the simulation model as no sufficiently precise estimate of their ferromagnetic properties could be obtained.

4.3.1 Changes to the simulation model based on surveys in the cavern

Surveys of the inner (top and bottom) planes of the yoke window, *i.e.* the pole faces, were performed on May 28th, 2021 [42]. Nine targets were placed at points along the pole face edges. Two (four) additional targets were positioned in the central region of the bottom (top) plane. The position of all targets with respect to the LHCb coordinate system was measured. The resulting sets of points for each face were used to define two planes by minimising the sum of the squared normal distances between the sets of measured coordinates and the plane. The fitted planes were compared to the planes spanned by the pole faces in the simulation model, and adjustments were made to the latter planes to bring them into coincidence with the fitted planes by rotating them around the origin of the coordinate system and shifting them in the y direction. These rotations and shifts, summarised in Table 4.1, were also propagated to the shims on the sides of the pole faces. The top and bottom blocks of the magnet yoke were then redrawn in the simulation model, shifting only the four corner points of each pole face, and

Table 4.1 – Transformations applied to the magnet pole faces in the OPERA simulation model.

	axis of rotation	rotation angle	shift in <i>y</i>
top face	(0.956, 0.124, 0.265)	0.72 mrad	3.5 mm
bottom face	(0.880, 0.117, -0.460)	0.18 mrad	-2.8 mm

keeping the other corners fixed. The shims were redrawn with all their corner points rotated and shifted according to the plane of the pole face they sit on.

Surveys of the magnet coils were taken to compare their positions between the “magnet on” and “magnet off” states in June 2012 [43] with photogrammetric retro-reflective targets placed on each of the five pancake triplets that make up a coil. The targets were pasted to both coils on the upstream side, with four targets per pancake triplet, two each on the A-side and C-side, near the bend of the coils into the downstream region inside the magnet. The targets were found to shift outwards by a few mm when the magnet was on, with respect to their positions when the magnet was off; an upstream shift in the *z* direction on the order of 5 mm was also observed. Consequently, the “opening” shifts in *x* and *y* were applied to those conductor blocks in the coil simulation models that make up the tapering sides, with one block on each end of that region “stretched” by having only one of its faces shifted, to ensure that the coil models remain unbroken. The average of all *z* shifts observed across the survey targets was applied to each of the coils as a whole.

4.3.2 Additional ferromagnetic structures in the simulation model

The first and most impactful ferromagnetic addition to the simulation model concerned the yoke support chariot. It was designed in OPERA with a simplified geometry, taking care that the volume and position of each part correspond to the values in a detailed CAD model of the LHCb spectrometer and its surroundings. With a total weight of over 70 tons, it has the largest impact on the magnetic field out of all the material changes made to the model, being essentially an extension of the yoke and bringing a sizeable up-down asymmetry into the model of the magnet environment. The description of the yoke support chariot in OPERA is shown in green in Figure 4.9.

The scaffolding on the A-side and C-side of the magnet is a collection of rather complex structures. While their combined mass is over 20 tons, they are mostly far enough from the magnet yoke that their description in OPERA can be simplified. For the simulation model, platforms and stairs were grouped into seven different components. Ten vertical support pillars of the platforms and stairs make up the rest of the scaffolding parts. For each part (stair and/or platform components as well as pillars), the description in OPERA was based on only three attributes:

- the (x, y, z) coordinates of the structure's centre of mass;
- the 3×3 inertia tensor with respect to the centre of mass;
- the total volume.

These attributes were calculated from the same CAD model of LHCb that provided the yoke support chariot geometry. Each of the 17 scaffolding parts is represented in the OPERA model by a rectangular iron block with the same centre of mass, moments of inertia, and volume as the original complex structure. The scaffolding parts are shown in olive in Figure 4.9.

The clamps that hold the coil in place, whose combined volume is a little less than 1 m^3 , are represented in the OPERA model by iron blocks of dimensions $25 \text{ cm} \times 16.5 \text{ cm} \times 80 \text{ cm}$ for the clamps along the straight sections of the coils and $30 \text{ cm} \times 25 \text{ cm} \times 60 \text{ cm}$ for the outermost clamps at the bent section, which were added to the magnet in the period between Runs 1 and 2. The clamps are sand-coloured in Figure 4.9.

4.4 Interpolation and extrapolation

In tracking and reconstruction at LHCb, the field values at a point defined by coordinates (x, y, z) need to be calculated from the map, which is defined on a cubic grid. In the LHCb software, this calculation uses linear interpolation from the vertices of the grid cube that contains (x, y, z) . For magnetic fields that vary slowly with the grid spacing, this approach gives accurate enough estimates of the fields. A more precise interpolation method that uses polynomials of order 2 in the distance from the cube vertices which intrinsically respect Maxwell's equations [44], was implemented in the course of this work for the computation of field values at the measurement points, for purposes of comparison between measurements and simulation data.

The simulated field values exported from OPERA have discontinuities at the material boundaries of the simulation model. Grid vertices that lie within material should therefore not be used to interpolate field values at points that lie outside of the material boundaries, *i.e.* in air. As the boundaries of the LHCb acceptance lie close to the material boundaries, especially within the window of the magnet yoke, a point in the acceptance may lie inside a grid cube that has at least one vertex inside some material in the model. Due to the shifts that were




Figure 4.10 – Comparison of field values at grid vertices in air (blue circles) and grid vertices in material with extrapolated field values (green circles) and the field values exported directly from the simulation (red crosses) at $x = -2$ m and $z = 4.5$ m for the three field components B_x (top left), B_y (top right), and B_z (bottom left) as well as for the magnitude of the field $|\vec{B}|$ (bottom right).

applied to LHCb field maps preceding this work (see Section 4.1), such discontinuities affected the field values in the maps even within the LHCb acceptance, at large positive values of y in the magnet yoke window. For the field maps produced in the course of this work, the grid vertices that lie within material boundaries have values assigned to them that are extrapolated, using the same parametrisation of the field in Maxwell-compliant polynomials as in the interpolation, from a nearby grid cube that is fully in air, *i.e.* that does not have any vertices in material. The main goal of this extrapolation is the improvement of the field description in the LHCb acceptance. The search algorithm for a suitable grid cube can be summarised as follows: Grid cubes near the original vertex that lies within material are checked until one cube fully in air is found. To determine the order in which nearby cubes are checked, the magnitudes of the x and y coordinates of the original vertex are compared. If $|x| \leq |y|$, the algorithm prioritises neighbouring cubes in the direction of lower values of $|y|$ (towards the beam line). Then, the neighbouring cubes in the direction of lower $|x|$ are checked. The next search direction is downstream towards higher values of $|z|$ (or higher values of z , since this coordinate is positive for all materials in the simulation model) due to the pyramidal shape of the LHCb acceptance. If $|x| > |y|$, the search order of these two directions is reversed. All cubes with a given distance to the original vertex are checked before widening the search to the next layer of surrounding cubes (nearest neighbours, then next-to-nearest neighbours, and so on). The search ends when a suitable cube that is fully in air is found, or when the distance to the original vertex reaches 5 m. In the latter case, the simulated field values of the original vertex are used as is. A comparison between the field values without this extrapolation algorithm (using the simulated values directly for vertices in material) and those obtained from extrapolation is shown in Figure 4.10. Unlike the field values from the simulation for grid points in material, those obtained from extrapolation give smooth continuations of the fields in air for the first few grid points inside the material, which ensures that the interpolation of the field values close to acceptance boundaries is not corrupted.

4.5 The magnetic field map for Run 3

The field map produced for use in LHCb during Run 3 is based on results of field simulations with the updated simulation model described in Section 4.3. The measurement data from 2021 is taken into account in a similar way as for the maps produced in 2011 and 2014, by shifting the grid coordinates to reduce the discrepancy with the measured field values. In this case, however, the map that is shifted comes out of the updated simulation, with the extrapolation procedure from Section 4.4 applied. The extrapolation does not influence the result of the fit to data, as the measurements lie well within the acceptance.

Figure 4.11 shows the difference between the measured B_y values for the “MagDown” polarity and the values interpolated from a simulated grid based on the updated simulation model. The analogous Figures A.1 and A.2 for the other two field components are available in Appendix A.1. The data-simulation differences are also shown as histograms in Figure 4.12. A comparison with Figures A.4, A.3, A.5, and A.6, which show the differences between the measurements and the magnetic field map used in LHCb in Run 2, indicates that the new simulation model alone improves the field description with respect to the previous map for B_x and B_z . The




Figure 4.11 – Differences between the corrected magnetic field measurement values for the B_y component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, before shifting the magnet.




Figure 4.12 – Histograms of the field differences (in Tesla) of Figures A.1, 4.11, A.2.

Table 4.2 – Transformations applied to the magnet in the OPERA simulation model based on a fit of the simulated map to field measurements from January 2021. The magnet is shifted and then rotated around the pivot point at (0, 0, 4.3158 m). The rotation is given in the form of extrinsic Euler angles around axes that are parallel to the LHCb coordinate system.

Translations [cm]				Rotations [mrad]			
x	y	z	total	x	y	z	total
-1.2548	0.1517	-1.6069	2.0444	-2.6559	1.7863	-0.0045	3.2011

z dependence of the sign of the difference that can be seen in Figure 4.11 is indicative of a relative shift in z .

To quantify the z translation of the field, as well as translations in x and y , and rotation parameters around a pivot point at (0, 0, 4.3158 m) (a point in the central region of the magnet that was also used in the production of the 2011 LHCb field map [37]), a minimisation fit is performed. The minimised quantity is the sum of the squared differences in the three B -field components between the data and the simulation, over the measurement points. The simulated field values at the measurement positions are interpolated from the grid using the Maxwell-compliant polynomials of degree 2 from Ref. [44]. The measurement coordinates are shifted in the fit iterations until the minimum is reached. The fit to the measurement data uses measured field values after the rotation corrections described in Section 4.2. The results for both magnet polarities are compatible, and the “MagDown” results are used in the final map. The opposite of the final shift to the measurement coordinates is applied to the magnet (the coils, current lines, yoke, shims, yoke support chariot, and coil support clamps) in the simulation model, and magnetic fields are re-computed for the new geometry, *i.e.* the simulation is re-run. Shifting the magnet within the simulation model instead of the whole field map changes the fields less in regions far from the pivot point of the rotation. The transformations applied to the magnet in the model are shown in Table 4.2.

Figures A.7 to A.10 in Appendix A.1 show that the differences between the measured and simulated field values are reduced after shifting the magnet model and that the systematic z dependence of the B_y difference sign seen in Figure 4.11 is compensated.

Global and local scale factors for the field values were also considered but not incorporated in the final map. Likewise, a local adjustment of the map in the measurement region, where the field values are high and/or vary significantly over short distances (which is where accuracy of the field map is paramount) was tested but not chosen for the final Run 3 map. The choice of map for use in Run 3 was based on the performance of the detector alignment and track reconstruction, which will be introduced in the following.

The precision and accuracy of the reconstruction of charged particle momenta directly influences the resolution and accuracy of the reconstructed invariant masses. For the decay $D^0 \rightarrow K^-\pi^+$, for example, the D^0 invariant mass can be computed from the reconstructed four-momenta of the K^- and the π^+ , charged long-lived particles, which leave signals (called “hits”) in the tracking detectors. The particles’ trajectories are reconstructed from the tracker hits and the four-momenta of the K^- and π^+ immediately after the D^0 decay are calculated in the reconstruction with input from the magnetic field map. If the momentum reconstruction is accurate, the distribution of the reconstructed D^0 invariant mass will be centred at the D^0 mass of $1864.84 \pm 0.05 \text{ MeV}/c^2$ [45]. If the momentum reconstruction is precise, the distribution of the reconstructed D^0 invariant mass will be narrow. A narrower mass peak means that the resolution is better.

The effect of the B field map quality on the momentum reconstruction is inextricably linked to that of the alignment performance. “Alignment” here refers to the fine-tuning of the position settings of the tracking detectors in the reconstruction framework (for details see e.g. Ref. [46]) so that the distance between the hit coordinates and the intersections of the reconstructed trajectories with the tracker planes is minimised. There is constant interplay, in the development (in terms of curvature) of the trajectories, between alignment, momentum calculation, and the magnetic field values in the (extrapolated) particle trajectories in this iterative procedure, so any measure of quality of the momentum reconstruction refers to both the magnetic field map and the alignment.

Mass peaks are shown in Figure 4.13 for the D^0 and J/ψ mesons, from their decays into $K^-\pi^+$ and $\mu^+\mu^-$, respectively. The peaks on the left-hand side have been calculated using the most recent previous version of the magnetic field map used in LHCb, with its associated alignment, and the right-hand side peaks are based on reconstruction using the new Run 3 map with its new alignment. All data shown in this figure and in Figure 4.14 were recorded by the LHCb experiment in 2024. The statistical uncertainties of the peak position and width are of the order of $0.2 \text{ MeV}/c^2$ for the J/ψ data sets and of the order of $0.05 \text{ MeV}/c^2$ for the D^0 data sets. The J/ψ peak in the new configuration is $2.4 \text{ MeV}/c^2$ closer to the known J/ψ mass [45] than in the old configuration. The resolution on the J/ψ mass does not change significantly. The D^0 peak in the new configuration is $1.66 \text{ MeV}/c^2$ closer to the known D^0 mass [45] than in the old




Figure 4.13 – The reconstructed J/ψ (top) and D^0 (bottom) mass in Run 3 data obtained with the Run 2 map (left) and the Run 3 map (right). The distributions are described with a linear background (blue) and a Gaussian signal peak (red). The green dashed lines indicate the known J/ψ and D^0 masses [45].

configuration. There is a slight (less than $1 \text{ MeV}/c^2$) improvement in the D^0 mass resolution. The improvements of the accuracy of the reconstructed mass and on the resolution obtained with the Run 3 map with respect to the Run 2 map are small but statistically significant.

The quality of track reconstruction and of the B field map may have spatial variations. One way to test for this is to check for a dependence of the reconstructed meson mass on its direction of flight. The slope parameters t_x and t_y are the ratios of the transverse momentum components p_x and p_y respectively to the forward momentum p_z . A reconstructed decay with $t_x > 0$ and $t_y > 0$ is likely to have both daughters' tracks in the first quadrant of the detector. Figure 4.14 shows the reconstructed J/ψ mass as a function of t_x and t_y . The reconstructed J/ψ mass shows a smaller dependence on t_y for the Run 3 map than for the Run 2 map.




Figure 4.14 – The reconstructed J/ψ mass as a function of the slope parameters t_x (top) and t_y (bottom) for the Run 2 map (left) and the Run 3 map (right). The blue dashed line indicates the true J/ψ mass [45].

Chapter 5

Data selection and B mass fit

5.1 Selection requirements

This work uses data recorded by the LHCb experiment in 2011, 2012, 2015, 2016, 2017, and 2018. Due to the high collision rate and consequent data volume, not every event can be retained, or even recorded. There are several stages of classification, interspersed with (partial) reconstruction, between the bunch collisions at the LHCb interaction point and the data available for physics analysis. The purpose of this process is to retain as much useful data as possible, while keeping the computational and memory loads manageable low.

The purpose of the LHCb trigger system is to identify potentially interesting events for retention, the triggered action being the recording of the event. Full reconstruction of all events is too expensive, so the triggering relies on simple markers such as high-energy deposits in some detectors that indicate a hard collision or an otherwise interesting event. For this work, at the hardware trigger level, events are selected that present a high-energy deposit in the ECAL. The high-level trigger requirements select B candidates that have one or more charged tracks in addition to a high transverse-energy photon candidate. For Run 1 data, a trigger line that reconstructs B candidates from three charged tracks, independently of the presence of a photon candidate, is also used. Table 5.1 summarises the trigger lines.

The data that pass at least one of these trigger lines at each level are loosely filtered in the “stripping” stage for topological correspondence with a $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ signal decay. Events with soft final-state particles are rejected, as are events with a low score for a classifier that discriminates photons from neutral hadrons (`gamma_CL`). Table 5.2 summarises the selection requirements at this stage.

The next batch of selection requirements is referred to as the “offline” selection. Like the stripping, its purpose is to select as many events with correctly reconstructed $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays as possible, while keeping the number of other events, referred to as background, rela-

Chapter 5. Data selection and B mass fit

Table 5.1 – Triggering requirements. An event must pass at least one of the lines listed at each level.

L0	L0Photon L0Electron	
	Run 1	Run 2
HLT1	Hlt1TrackAllL0	Hlt1TrackMVA
	Hlt1TrackPhoton	Hlt1TwoTrackMVA
HLT2	2011	2012
	Hlt2RadiativeTopoPhotonL0 Hlt2RadiativeTopoTrackTOS Hlt2Topo3BodyBBDT	Hlt2RadiativeTopoPhoton Hlt2RadiativeTopoTrackTOS Hlt2Topo3BodyBBDT

Table 5.2 – Requirements applied for the stripping selection. The detachment of the trajectory of a particle from the PV is quantified by the impact parameter, which is the minimal distance from the extrapolated trajectory to the PV; the related figure of merit χ_{IP}^2 should be small for the B candidate and large for the tracks of the final state particles, which do not come from the PV. The B DIRA is the cosine of the angle between the B direction and the line connecting the PV and the B decay vertex.

Variable	2011–2012	2015–2018	Unit
Track transverse momentum p_{T}	> 300	> 300	MeV/c
Track momentum p	> 1000	> 1000	MeV/c
Track fit quality χ^2/ndf	< 3	< 3	
Track detachment χ_{IP}^2	> 16	> 20	
Track ghost probability	< 0.4	< 0.4	
Tri-track $\sum p_{\text{T}}$	> 1500	> 1000	MeV/c
Tri-track vertex fit quality χ_{vtx}^2	< 10	< 9	
Tri-track vertex separation from PV χ_{VS}^2	> 0	> 0	
Tri-track invariant mass	$\in [0, 7900]$	$\in [0, 7900]$	MeV/c^2
Photon E_{T}	> 2000	> 2000	MeV
γ confidence level (gamma_CL)	> 0	> 0	
Photon and tracks $\sum p_{\text{T}}$	> 5000	> 3000	MeV/c
B DIRA	> 0	> 0	
B vertex fit quality $\chi_{\text{vtx}}^2/\text{ndf}$	< 9	< 9	
B detachment χ_{IP}^2	< 9	< 9	
B candidate mass	$\in [2900, 9000]$	$\in [2400, 6500]$	MeV/c^2

tively low. A b hadron may decay into several final-state particles of which three are charged and one is neutral. If the reconstruction, looking for $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays, picks up some but not all of the final-state daughters of a true decay, the event contributes to what is called the partially reconstructed background. Some of the topological selection requirements at the stripping stage are intended to reduce this background. Another similar criterion is defined in the offline selection: each of the tracks in the event that is not associated to one of the final-state particles in the putative $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decay is merged in turn with those four “signal” finals, and a decay vertex is reconstructed from the four charged tracks and the neutral

candidate. If the candidate is actually from a signal decay, and if the vertex is spatially isolated from spurious tracks from the rest of the event, then adding any other track to the vertex decreases the vertex quality ($\Delta\chi^2 > 0$). Therefore it is required that there be a significant drop ($\Delta\chi^2 > 5$) in vertex quality when any track that is not part of the candidate is added to the decay vertex. If this is not the case, the candidate is a partially reconstructed decay, or is not well enough isolated in space to be discriminated from such a background.

There are background processes, some partially reconstructed among them, which peak in some invariant mass distributions. These “peaking” backgrounds can be reduced by requiring that an invariant mass combination, for example the invariant mass of the $K\pi\pi$ system, lie outside of the region where the background peaks. The process $B^+ \rightarrow \bar{D}^0 \rho^+$, where the \bar{D}^0 decays to $K^+ \pi^- \pi^0$ and the ρ^+ to $\pi^+ \pi^0$, can mimic a $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ signature when one of the final-state π^0 is lost, and the other one is misidentified as a photon. In case the π^0 from the \bar{D}^0 decay is reconstructed as a photon, and the π^0 from the ρ^+ decay is lost, the invariant mass of the K^+ , the π^- , and the photon candidate (assigning the π^0 mass to it), is close to the D^0 mass. In the other case, where the π^0 from the ρ^+ is reconstructed as a photon and the π^0 from the \bar{D}^0 is lost, the invariant mass of the $\pi^+ \pi^0$ system is close to the ρ^+ mass. By requiring the invariant masses concerned to be a good deal higher than the D^0 mass, respectively the ρ^+ mass, the contributions from this background in the data are significantly reduced.

Finally, each of the particles in the final state has one dangerous doppelganger that can be mistaken for it. To reduce background from these events with misidentified particles, the particle identification (PID) capabilities of the LHCb detector and software are exploited. For example, the dangerous doppelganger of the K^+ is the π^+ (and vice versa). There is a multivariate classifier that uses information from the PID sub-systems to determine the probability that the K^+ candidate is a K^+ , and another classifier that yields the probability that it is actually a π^+ . The output of these two classifiers is combined into one quantity by multiplying the K^+ probability with one minus the π^+ probability. An event is rejected if this combined variable has a value less than 0.2. The π^\pm candidates have the inverted requirement: their pion probability, times one minus their kaon probability, must be greater than 0.2. These PID classifiers do not perform well for very high momenta, so the high- p tails of the charged final-state particle candidates are discarded by requiring that $p < 100 \text{ GeV}/c$ for the K^+ , π^+ , and π^- candidates. The photon’s doppelganger is the neutral pion π^0 . Since the photon is not charged, photon identification can only depend on the information from the calorimeters, specifically the ECAL. A high-energy photon can transform into a positron-electron-pair through interaction with the calorimeter material. The positron and electron interact further, mostly through emission of Bremsstrahlung photons, which may convert into more $e^+ e^-$ pairs producing more photons and electrons in turn, until a whole shower of electromagnetic particles is detected and its energy is measured. The ECAL is segmented, so some information about the distribution of the shower (energy) in space is available. This becomes important when it comes to differentiating an ECAL shower that was caused by a photon from one caused by a π^0 after decaying into a pair of photons. For relatively low-energy pions, the momentum of the photons perpendicular to the initial π^0 direction is large enough compared

Table 5.3 – Requirements applied for the offline preselection, excluding the multivariate classifier for rejection of combinatorial background.

Variable	2011–2012	2015–2018	Unit
Photon E_T	> 3000	> 3000	MeV
Photon E_T	< 120000	< 120000	MeV
K^+ momentum	< 100000	< 100000	MeV/ c
π^+ momentum	< 100000	< 100000	MeV/ c
π^- momentum	< 100000	< 100000	MeV/ c
B meson p_T	> 5000	> 4000	MeV/ c
B meson isolation $\Delta\chi^2$	> 5	> 5	
γ/π^0 separation (IsPhoton)	> 0.9	> 0.9	
γ confidence level (gamma_CL)	> 0.2	> 0.2	
ProbNN $_K(K^+) \times (1 - \text{ProbNN}_\pi(K^+))$	> 0.2	> 0.2	
ProbNN $_\pi(\pi^+) \times (1 - \text{ProbNN}_K(\pi^+))$	> 0.2	> 0.2	
ProbNN $_\pi(\pi^-) \times (1 - \text{ProbNN}_K(\pi^-))$	> 0.2	> 0.2	
$K^+\pi^-\pi^+$ mass	$\in [1100, 1900]$	$\in [1100, 1900]$	MeV/ c^2
$K^+\pi^-\pi^0$ mass	> 2200	> 2200	MeV/ c^2
$\pi^+\pi^0$ mass	> 1100	> 1100	MeV/ c^2
$\pi^-\pi^0$ mass	> 1100	> 1100	MeV/ c^2

to their parallel momentum that they acquire an appreciable separation before reaching the ECAL. When this separation is large enough, their showers can be “resolved” (reconstructed separately). The two photons from a high-energy π^0 decay, on the other hand, reach the ECAL so close together that their showers overlap. This is referred to as a “merged” π^0 , as opposed to the resolved case at lower energy. A merged shower can look like the shower from a true single photon with an energy close to the energy of the initial π^0 . The best discrimination power has been achieved using a multivariate classifier with several variables that quantify the shower shape (size, concentration, asymmetry, etc.) [47]. The resulting discrimination score is called IsPhoton and is close to 1 when the photon confidence is high, and close to 0 when the shower resembles that of a π^0 . It is required to be higher than 0.9 for this analysis. Because the classifier does not distinguish neutral candidates with very high energies well, the transverse energy of the photon candidate is required to be below 120 GeV. All the offline selection requirements defined up to this point are summarised in Table 5.3.

The requirements on the PID variables are some of the strictest. Their efficiencies, which are not reproduced well in MC, are evaluated on data samples prepared for this purpose, where any background is corrected for. This is done using PIDCalib2 [48] for the ProbNN variables, and GammaPi0SeparationCalib [49] for the IsPhoton variable. The calibration samples are binned in the variables p and η for the charged, and p_T and η for the neutral particles. Each MC event is assigned a weight equal to the product of the efficiency of the PID requirements on data in its kinematic bin, instead of applying the PID requirements on MC.

For all quantities that can be defined and calculated to characterise a pp collision event, there is some overlap between the distribution in a randomly selected sample of all recorded events and the distribution specific to a signal-only sample. A decay candidate that is not signal and does not belong to any of the previously mentioned background categories (misID, partially reconstructed), exhibiting a signal-like signature by simple stochastic accident, is termed “combinatorial” background, being the result of a random combination of tracks, energy deposits, or noise. The generic nature of this background, contrasted with the specificity of a signal signature, can be used to reject some of these combinatorial events. However, simple one-dimensional rejection criteria are not optimal for this background type. Even with very overlapping variable distributions, though, a multi-dimensional approach can squeeze some discrimination power out of the parameter space. In this work, this is done using a multivariate classifier with boosted decision trees (BDT). The classification is trained on signal MC (“norm”, see Section 6.3) as a signal proxy. The selection efficiency for combinatorial events is so low that using MC for the background proxy would be computationally expensive to an unreasonable extent. Instead, the background proxy is taken from events in data. Events with a B candidate mass far away from the signal peak are background events. In the low mass range, these can be partially reconstructed backgrounds, but in the high mass range, only combinatorial events are present. The training sample is chosen from data where all the previously described selection steps have been applied and the B candidate mass is greater than 5620 MeV/ c^2 . The signal and background proxy samples differ by construction in their B mass range. So any variable used in the classifier must be independent of the B mass. As one sample is from MC and the other one from data, these variables must also be chosen among those which are well reproduced in MC. The list of variables chosen is given in Table 5.4. Each variable with CHI2 or Chi2 in its name characterises the quality of the reconstruction of some part of the event under the assumption that the event is a $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ signal event; generally, the better the match, the lower the value of the variable. Exceptions to that rule are Kplus_MINIPCHI2, piplus_MINIPCHI2 and piminus_MINIPCHI2, which are related to the minimum distance of the trajectory to the primary vertex (PV), which is the proton-proton interaction point of origin of the B candidate [50]. The charged hadrons should come from the B decay vertex, which is displaced from the PV. K_res_ENDVERTEX_CHI2 is the χ^2/ndf of the $K^+ \pi^- \pi^+$ vertex. The smaller variable B_IPCHI2_OWNPV is, the more likely the PV is to be the origin of the B candidate. B_DIRA_OWNPV is also related to the compatibility of the B candidate and the PV, but rather than a fit χ^2 , it is the cosine of the angle between the B momentum and the line between the PV and the B decay vertex. B_OWNPV_CHI2 is the χ^2/ndf of the PV itself. The greater the flight distance between the PV and the B decay vertex, the higher the value of B_FDCHI2_OWNPV. This variable is another exception of the general rule that a smaller χ^2 variable is more signal-like. Another exception is B_SmallestDeltaChi2OneTrack, which is small when there is a track in the event not belonging to the signal candidate that is compatible with the B decay vertex, and the last χ^2 variable that should be large for signal events is B_MINIPCHI2NEXTBEST, which characterises the compatibility of the B candidate to originate at another proton-proton interaction point (the next best candidate for the primary vertex), and is only available for Run 2 data. Also available in Run 2 only are the cone iso-

Table 5.4 – Variables used in the BDT to reject combinatorial background. The Run 2 variables `K_res_ENDVERTEX_CHI2` and `B_VTXISODCHI2ONETRACK` are aliases of and equivalent to the Run 1 variables `K_1_1270_plus_ENDVERTEX_CHI2` and `B_SmallestDeltaChi2OneTrack` respectively.

2011–2012 (Run 1)	2015–2018 (Run 2)
<code>B_FDCHI2_OWNPV</code>	
<code>K_1_1270_plus_ENDVERTEX_CHI2</code>	<code>K_res_ENDVERTEX_CHI2</code>
<code>B_SmallestDeltaChi2OneTrack</code>	<code>B_VTXISODCHI2ONETRACK</code>
<code>Kplus_MINIPCHI2</code>	<code>Kplus_MINIPCHI2</code>
<code>piplus_MINIPCHI2</code>	<code>piplus_MINIPCHI2</code>
<code>piminus_MINIPCHI2</code>	<code>piminus_MINIPCHI2</code>
<code>B_IPCHI2_OWNPV</code>	<code>B_IPCHI2_OWNPV</code>
<code>B_DIRA_OWNPV</code>	<code>B_DIRA_OWNPV</code>
<code>B_OWNPV_CHI2</code>	<code>B_OWNPV_CHI2</code>
	<code>B_MINIPCHI2NEXTBEST</code>
	<code>B_B_CONEP_1.0</code>
	<code>B_B_CONEPTASYM_1.0</code>
	<code>B_B_CONEPTASYM_1.35</code>
	<code>B_B_CONEPTASYM_1.7</code>
	<code>B_Gamma_CONEPTASYM_1.0</code>
	<code>B_Gamma_CONEP_1.0</code>

lation variables which give the sum of the absolute momenta of tracks in a cone around the B trajectory (`B_B_CONEP_1.0`) or the photon trajectory (`B_Gamma_CONEP_1.0`), as well as the p_T asymmetry, which is the normalised difference between the p_T of the B candidate (`B_B_CONEPTASYM_1.0`, `B_B_CONEPTASYM_1.35`, `B_B_CONEPTASYM_1.7`) or the photon candidate (`B_Gamma_CONEPTASYM_1.0`) and the sum of the p_T of all the tracks in the cone. The numbers in the name of the cone isolation variables give the radius of the cone, which is calculated from the pseudorapidity η and the azimuthal angle ϕ (in radians) as $\sqrt{(\delta\eta)^2 + (\delta\phi)^2}$.

Since there are operational differences between data-taking years, and the sample sizes are sufficiently large, each year gets its own classifier. In fact, for the purposes of simultaneous implementation and testing, eight classifiers are trained for each year on seven eighths of the signal and background proxy samples, with the remaining eighth used to test and quantify the classifier’s performance. So as not to bias the estimation of the efficiency of the cut on the classifier output, the data and MC events are evaluated by the classifier for which they were in the testing, and not in the training, sample.

With any suboptimal classifier, there is a trade-off between the false positive rate (the proportion of true background events classified as signal) and the false negative rate (the proportion of true signal events classified as background). A stricter requirement on the classifier output score lowers the false positive rate, but also rejects more signal events. Conversely, a loose requirement leads to a lower loss of signal, *i.e.* a higher true positive rate, at the cost of higher background contamination. Plotting the false positive rate vs. the true positive rate leads to a

5.1. Selection requirements

Table 5.5 – Optimised requirements on the multivariate (BDT) classifier scores to reject combinatorial background events and corresponding AUC.

Year	Requirement	AUC
2011	> 0.45	0.946 ± 0.010
2012	> 0.60	0.940 ± 0.006
2015	> 0.60	0.948 ± 0.013
2016	> 0.55	0.961 ± 0.003
2017	> 0.55	0.957 ± 0.004
2018	> 0.45	0.961 ± 0.003

Table 5.6 – Fiducial requirements on invariant masses of the hadronic systems.

Hadronic system	Low limit	High limit
$K^+ \pi^-$	$650 \text{ MeV}/c^2$	$1200 \text{ MeV}/c^2$
$\pi^+ \pi^-$	$530 \text{ MeV}/c^2$	$900 \text{ MeV}/c^2$
$K^+ \pi^- \pi^+$	$1100 \text{ MeV}/c^2$	$1800 \text{ MeV}/c^2$

curve called the ROC-curve, for *receiver operating characteristic*. The BDT performance can be characterised by the area under the ROC-curve (AUC) [51]. An AUC of 0.5 means that the BDT performs no better than a random selection. An AUC of 1 is a perfect score. The cut on the BDT output is chosen, for each year, so that it maximises the product of the signal significance and the signal purity, $\frac{S}{\sqrt{S+B}} \times \frac{S}{S+B}$ where S and B are respectively the expected numbers of signal and background events. These numbers are computed from a fit to the B candidate mass distribution in the data sample with the offline selection applied. The fit components are described in Section 5.2. Events with B candidate masses between $4300 \text{ MeV}/c^2$ and $6500 \text{ MeV}/c^2$ are included in the fit. The fitted signal shape and combinatorial background shape are integrated in the B candidate mass range between $5100 \text{ MeV}/c^2$ and $5500 \text{ MeV}/c^2$ to obtain the values for S and B . The optimised requirements for the different years, as well as the AUC, are given in Table 5.5. The ROC-curves for one of the classifiers for each year out of the eight classifiers per year are shown in Figure 5.1.

A final set of offline selection requirements is related to the performance of the amplitude fit. One of the challenges of this analysis is the large number of decay chains of B^+ mesons that lead to the $K^+ \pi^+ \pi^- \gamma$ final state. This multiplicity is due to the many intermediate resonances in the hadronic system (whereas intermediate resonances decaying to one or two hadrons and a photon are suppressed so as to be negligible by requiring a very high-energy photon). By paring down the list of decay chains that must be considered in the amplitude analysis, the model building described in Section 6.4 can be rendered more robust. The intermediate resonances show up in different regions of the phase-space of invariant masses of the hadronic system, depending on their mass peak and width. The $m(K^+ \pi^-)$ spectrum is dominated by the $K^*(892)^0$, while the main peak of the $m(\pi^+ \pi^-)$ spectrum is that of the $\rho(770)^0$. The importance of other resonances, such as the $K^*(1430)$ or the $f_2(1270)$, can be reduced by requiring the

Chapter 5. Data selection and B mass fit




Figure 5.1 – Performance of the multivariate classifier to reject combinatorial background: for each data-taking year, the ROC-curve of one of the eight classifiers trained for that year is shown. The true positive rate and false positive rate obtained with the optimised threshold for that year is indicated in red. For the mean of the AUC for all eight classifiers of each year, see Table 5.5.

5.1. Selection requirements

invariant masses of the $K^+\pi^-$ and $\pi^+\pi^-$ systems to lie within the bulk of the $K^*(892)^0$ and $\rho(770)^0$ peaks, respectively. Similarly, the description of the high-mass tail of the $m(K^+\pi^+\pi^-)$ spectrum can be simplified by constraining the allowed values for that invariant mass. These “fiducial” requirements are listed in Table 5.6. They are not included in the list of offline selection requirements of Table 5.3 because they were not applied on the samples that were used to train the multivariate classifier for combinatorial background rejection. The numbers of events after offline selection and fiducial requirements are 3 041 in 2011, 8533 in 2012, 2 525 in 2015, 20 623 in 2016, 19 911 in 2017, and 25 116 in 2018, for a total of 79 749.

5.2 Fit to B candidate invariant mass

Information on the make-up of the data sample after the selection has been applied can be obtained from the B candidate mass distribution. It shows a clear signal peak at the B^+ mass. By describing this peak mathematically, the total number of signal events in the sample can be calculated. Different background contributions to the data sample have different B mass distributions, which can be characterised and quantified using MC simulations as well as the data sample itself. This description of the B mass distribution in terms of the shapes and yields of the signal and the significant backgrounds allows the subtraction of the backgrounds from the data, using a weighting procedure that will be introduced in Section 5.3. This background-subtracted data sample is used in the amplitude fit in Chapter 6. In the B mass fit, data from a wide mass range between 4300 MeV/ c^2 and 6500 MeV/ c^2 is used. Once the descriptions of the shapes and the yields are fixed, the B mass range can be constrained to 5000 – 6500 MeV/ c^2 for the background subtraction. The data in the range 4300 – 5000 MeV/ c^2 is almost free from signal, but is retained for the B mass fit so that the shapes of the backgrounds that populate this region can be characterised.

The B mass distribution of the signal is modelled with a double-sided Crystal Ball (CB) probability density function (PDF) [52]:

$$f_{\text{CB}}(x; \mu, \sigma, \alpha_L, \alpha_R, n_L, n_R) = \mathcal{N} \times \begin{cases} A_L \left(B_L - \frac{x-\mu}{\sigma} \right)^{-n_L} & \text{for } \frac{x-\mu}{\sigma} \leq -|\alpha_L| \\ A_R \left(B_R - \frac{x-\mu}{\sigma} \right)^{-n_R} & \text{for } \frac{x-\mu}{\sigma} \geq |\alpha_R| \\ \exp\left(-\frac{(x-\mu)^2}{2\sigma^2}\right) & \text{for } -|\alpha_L| < \frac{x-\mu}{\sigma} < |\alpha_R| \end{cases}, \quad (5.1)$$

where \mathcal{N} is a normalisation factor and

$$\begin{aligned} A_L &= \left(\frac{n_L}{|\alpha_L|} \right)^{n_L} \exp\left(-\frac{\alpha_L^2}{2}\right), & B_L &= \frac{n_L}{|\alpha_L|} - |\alpha_L|, \\ A_R &= \left(-\frac{n_R}{|\alpha_R|} \right)^{n_R} \exp\left(-\frac{\alpha_R^2}{2}\right), & B_R &= -\frac{n_R}{|\alpha_R|} + |\alpha_R|. \end{aligned}$$

This smooth PDF is a Gaussian in the bulk of the peak, but has more degrees of freedom for the shapes of the tails on both sides. For each data-taking year, the tail parameters $\alpha_L, \alpha_R, n_L, n_R$ are determined from fits to the post-selection “norm” MC samples (see Section 6.3), with the PID efficiency weights applied. The mean μ and the width σ of the bulk of the peak, which is equal for both sides, are determined from the full B mass fit on the data distribution.

For the backgrounds from specific B meson decays other than the signal mode, the expected contamination, which is defined as the ratio of the number of background events to the number of signal events in a given sample, can be calculated from the branching fractions of the signal and the background decays and from their respective selection efficiencies. The branching fractions are fixed or estimated from existing measurements, while the selection efficiencies are determined by applying the selection detailed in Section 5.1 on simulated

5.2. Fit to B candidate invariant mass

Table 5.7 – List of samples of fully simulated events used to estimate the background contaminations. The $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ signal samples have been generated with the “val” model given in Table 6.3.

	Year	Number of events	Year	Number of events
$B^0 \rightarrow K^{*0} (\rightarrow K^+ \pi^-) \gamma$	2012	6.053×10^6	2016	1.013×10^6
$B^+ \rightarrow K^{*0} (\rightarrow K^+ \pi^-) \pi^+ \gamma$	2012	1.523×10^6	2016	2.012×10^6
$B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \pi^+$	2012	2.091×10^6	2017	1.016×10^6
$B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \rho^+ (\rightarrow \pi^+ \pi^0)$	2012	2.119×10^6	2016	1.007×10^6
$B^+ \rightarrow \bar{D}^{*0} (\rightarrow \bar{D}^0 (\rightarrow K^+ \pi^-) \gamma) \pi^+$	2012	1.032×10^6	2016	1.317×10^6
$B^+ \rightarrow \bar{D}^{*0} (\rightarrow \bar{D}^0 (\rightarrow K^+ \pi^-) \pi^0) \pi^+$	2012	1.008×10^6	2016	2.603×10^6
$B^+ \rightarrow K^{*+} (\rightarrow K^+ \pi^0) \pi^+ \pi^-$	2012	2.012×10^6	2016	1.034×10^6
$B^0 \rightarrow K_1(1270)^0 (\rightarrow K^+ \pi^- \pi^0) \gamma$	2012	1.127×10^6	2016	1.013×10^6
$B^+ \rightarrow K^+ \pi^+ \pi^- \eta (\rightarrow \gamma \gamma)$	2012	1.522×10^6	2016	0.514×10^6
$B^+ \rightarrow K_1(1270)^0 (\rightarrow K^+ \pi^+ \pi^-) \pi^0$	2012	0.621×10^6	2016	2.529×10^6
$B^+ \rightarrow a_1^+ (\rightarrow \pi^+ \pi^- \pi^+) \gamma$	2012	1.008×10^6	2016	0.511×10^6
$B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ (signal)	2012	0.232×10^6	2016	0.301×10^6

samples. The running conditions during Run 1 and Run 2 of the LHC were similar enough that it is sufficient to analyse simulation data for one data-taking year per run and assign the resulting contaminations and models to all the data-taking years of that run. For Run 1, the year chosen for simulation is 2012 and, for Run 2, it is 2016, except for the $B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \pi^+$ decay where a larger simulated sample was available for 2017 data-taking conditions. Also, studies which concern only the signal channel were often performed using the “norm” MC samples (see Section 6.3), of which one is available per data-taking year. The simulated MC samples used to estimate the background contaminations are shown in Table 5.7. The branching fractions, efficiencies, and expected contaminations of specific non-signal B meson decays that may significantly contribute to the post-selection data sample in the B mass fit range of $4300 - 6500$ MeV/ c^2 are shown in Table 5.8.

The proportion of double-mis-ID signal events, where the signal kaon is labelled as a pion and the pion of the same charge as the kaon is taken for the kaon, is reduced to the sub-permill level by the PID selection requirements and considered to be negligible.

For $B^+ \rightarrow a_1^+ (\rightarrow \pi^+ \pi^- \pi^+) \gamma$ decays, which can be misidentified as signal when one of the same-charge pions is taken for a kaon, the branching fraction is suppressed with respect to $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays by a factor of $\left| \frac{V_{td}}{V_{ts}} \right|^2 = 0.04$, where V_{td} and V_{ts} are the leading CKM-matrix elements for $b \rightarrow d$ (as in $B^+ \rightarrow a_1^+ (\rightarrow \pi^+ \pi^- \pi^+) \gamma$) and $b \rightarrow s$ (as in $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$) transitions, respectively. The efficiency of the selection on $B^+ \rightarrow a_1^+ (\rightarrow \pi^+ \pi^- \pi^+) \gamma$ decays is low enough to render their contamination negligible.

The background from $B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \rho^+ (\rightarrow \pi^+ \pi^0)$ decays is suppressed by the selection requirements placed on invariant mass combinations (see Section 5.1), but due to the relatively

Chapter 5. Data selection and B mass fit

Table 5.8 – List of potential peaking backgrounds and the quantities used to calculate their expected contaminations $C = \frac{\mathcal{B} \epsilon_{\text{sel}} \epsilon_{\text{gen}}}{\mathcal{B}_{\text{sig}} \epsilon_{\text{sel, sig}} \epsilon_{\text{gen, sig}}}$, in the B mass range $4300 - 6500 \text{ MeV}/c^2$ in Run 1 and Run 2 MC. Here, \mathcal{B} denotes the branching fractions, ϵ_{sel} the selection efficiencies, and ϵ_{gen} the generator-level efficiencies. The branching fractions are taken from Ref. [53] or calculated as indicated. The last line refers to the signal decay, with respect to which the contaminations are computed. When less than 3 candidates pass the selection, a 90% confidence level upper limit is quoted.

Decay chain	\mathcal{B}	Year	ϵ_{sel}	ϵ_{gen}	C
$B^+ \rightarrow a_1^+ (\rightarrow \pi^+ \pi^- \pi^+) \gamma$	$\frac{ V_{td} ^2}{ V_{ts} ^2} \mathcal{B}_{\text{sig}}$ $= (1.03 \pm 0.06) \times 10^{-6}$	2012 2016	$(4.05 \pm 0.74) \times 10^{-5}$ $(9.3 \pm 1.6) \times 10^{-5}$	0.20996 ± 0.00039 0.2150 ± 0.0085	$(2.34 \pm 0.43) \times 10^{-3}$ $(2.74 \pm 0.47) \times 10^{-3}$
$B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \pi^+$	$(6.6 \pm 0.3) \times 10^{-4}$	2012 2017	$< 1.43 \times 10^{-6}$ $< 2.95 \times 10^{-6}$	0.15843 ± 0.00037 0.16520 ± 0.00041	$< 4.02 \times 10^{-2}$ $< 4.30 \times 10^{-2}$
$B^+ \rightarrow K^{*+} (\rightarrow K^+ \pi^0) \pi^+ \pi^-$	$(2.5 \pm 0.3) \times 10^{-5}$	2012 2016	$< 1.49 \times 10^{-6}$ $< 2.92 \times 10^{-6}$	0.1608 ± 0.0066 0.14919 ± 0.00026	$< 1.59 \times 10^{-3}$ $< 1.44 \times 10^{-3}$
$B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \rho^+ (\rightarrow \pi^+ \pi^0)$	$(1.9 \pm 0.3) \times 10^{-3}$	2012 2016	$< 2.86 \times 10^{-6}$ $< 3.05 \times 10^{-6}$	0.14432 ± 0.00034 0.15172 ± 0.00024	$< 2.09 \times 10^{-1}$ $< 1.17 \times 10^{-1}$
$B^+ \rightarrow \bar{D}^{*0} (\rightarrow \bar{D}^0 (\rightarrow K^+ \pi^-) \pi^0) \pi^+$	$(1.23 \pm 0.05) \times 10^{-4}$	2012 2016	$< 2.97 \times 10^{-6}$ $< 1.16 \times 10^{-6}$	0.15252 ± 0.00030 0.15930 ± 0.00044	$< 1.49 \times 10^{-2}$ $< 3.04 \times 10^{-3}$
$B^+ \rightarrow \bar{D}^{*0} (\rightarrow \bar{D}^0 (\rightarrow K^+ \pi^-) \gamma) \pi^+$	$(6.7 \pm 0.3) \times 10^{-5}$	2012 2016	$< 2.90 \times 10^{-6}$ $< 2.28 \times 10^{-6}$	0.17229 ± 0.00034 0.17881 ± 0.00046	$< 8.96 \times 10^{-3}$ $< 3.63 \times 10^{-3}$
$B^0 \rightarrow K^+ \pi^- \pi^0 \gamma$	$(4.1 \pm 0.4) \times 10^{-5}$	2012 2016	$(3.1 \pm 2.7) \times 10^{-6}$ $(1.10 \pm 0.44) \times 10^{-5}$	0.20405 ± 0.00034 0.21559 ± 0.00047	$(6.8 \pm 6.0) \times 10^{-3}$ $(1.29 \pm 0.53) \times 10^{-2}$
$B^+ \rightarrow K^+ \pi^+ \pi^- \eta (\rightarrow \gamma\gamma)$	$\mathcal{B}_{\text{sig}} \frac{\mathcal{B}(B^+ \rightarrow K^{*+} \eta)}{\mathcal{B}(B^+ \rightarrow K^{*+} \gamma)} \mathcal{B}(\eta \rightarrow \gamma\gamma)$ $= (5.01 \pm 0.58) \times 10^{-6}$	2012 2016	$(3.47 \pm 0.16) \times 10^{-4}$ $(7.23 \pm 0.39) \times 10^{-4}$	0.16691 ± 0.00044 0.17311 ± 0.00040	$(7.71 \pm 0.85) \times 10^{-2}$ $(8.31 \pm 0.95) \times 10^{-2}$
$B^+ \rightarrow K^+ \pi^+ \pi^- \pi^0$	$\mathcal{B}_{\text{sig}} \frac{\mathcal{B}(B^+ \rightarrow K^{*+} \pi^0)}{\mathcal{B}(B^+ \rightarrow K^{*+} \gamma)}$ $= (4.48 \pm 0.69) \times 10^{-6}$	2012 2016	$(6.48 \pm 0.44) \times 10^{-5}$ $(1.358 \pm 0.077) \times 10^{-4}$	0.16960 ± 0.00026 0.18052 ± 0.00044	$(1.31 \pm 0.21) \times 10^{-2}$ $(1.46 \pm 0.23) \times 10^{-2}$
$B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ (signal)	$(2.58 \pm 0.15) \times 10^{-5}$	2012 2016	$(3.681 \pm 0.029) \times 10^{-3}$ $(6.592 \pm 0.043) \times 10^{-3}$	$(3.954 \pm 0.029) \times 10^{-2}$ $(4.429 \pm 0.033) \times 10^{-2}$	1 1

high branching fraction of this decay compared to that of the signal mode, they still contribute significantly to the partially reconstructed backgrounds in the post-selection data sample. The MC samples used to determine the selection efficiency on $B^+ \rightarrow \bar{D}^0 (\rightarrow K^+ \pi^- \pi^0) \rho^+ (\rightarrow \pi^+ \pi^0)$ events initially contain about 2M (1M) events for Run 1 (Run 2), see Table 5.7. The selection eliminates so many of these events, however, that the effective size of the post-selection samples is less than 1. This is not nearly enough to model the shape of this background species in the B candidate mass. It is more useful to introduce generic models for this and other partially reconstructed backgrounds which are characterised by the loss in reconstruction of one or two pions from the initial B meson decay.

A background that contaminates the signal region comes from $B^+ \rightarrow K^+ \pi^+ \pi^- \eta$ decays where the η decays into two photons, one of which is used to build the signal candidate, and the other of which is lost. The ratio of this decay's branching fraction to that of the signal decay is estimated from the ratio of branching fractions of the similar $B^+ \rightarrow K^{*+} \eta$ and $B^+ \rightarrow K^{*+} \gamma$ decays, which is equal to 0.49 ± 0.05 [45]. The branching fraction for $\eta \rightarrow \gamma\gamma$ decays of $(39.41 \pm 0.20)\%$ must be multiplied with this, and from the selection and generator-level efficiencies for the signal and for this background, we obtain expected contaminations from $B^+ \rightarrow K^+ \pi^+ \pi^- \eta (\rightarrow \gamma\gamma)$ decays of $(7.71 \pm 0.85)\%$ and $(8.31 \pm 0.95)\%$ of the number of signal events for Runs 1 and 2, respectively. The shape of this background in the B candidate mass is modelled with an Argus




Figure 5.2 – Mass distribution of simulated $B^+ \rightarrow K^+ \pi^+ \pi^- \eta$ events reconstructed as $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ in 2012 MC (left) and 2016 MC (right) with the result of a fit to an Argus PDF convolved with a double-tail Crystal Ball resolution function. The differences (residuals) between the numbers of bin entries and the average of the fitted blue curve over each bin in units of the uncertainty of the residuals are shown in the bottom plots, with grey bars indicating the deviation from zero and red dashed lines at ± 2 . No residuals are plotted for empty bins.

PDF [54] f_{Argus} , a standard function for describing partially reconstructed background mass distributions [55, 56],

$$f_{\text{Argus}}(x; m_0, c, p) = \mathcal{N} x \left(1 - \frac{x^2}{m_0^2}\right)^p e^{c\left(1 - \frac{x^2}{m_0^2}\right)}, \quad (5.2)$$

convolved with a function that accounts for the mass resolution, which is dominated by the resolution on the photon energy. The curvature c and power p of the Argus function are determined from a fit to the B mass distribution of simulated $B^+ \rightarrow K^+ \pi^+ \pi^- \eta$ events reconstructed as signal and with the full selection applied. In this fit to MC, the resolution function is fixed to the distribution of signal events in the “norm” MC. The end-point m_0 of the Argus PDF is fixed to the mean of that signal distribution. The normalisation \mathcal{N} of the Argus PDF is computed numerically in the fit range. The results of the fits for Run 1 and Run 2 can be seen in Figure 5.2. In the fit to the B mass distribution in data, the number of events for this background is pegged to the number of signal events multiplied by the contamination values quoted above. The curvature and power of the Argus are fixed to the results of the fit on MC shown in Figure 5.2, and the end-point is pegged to the signal mean. The resolution function is also pegged to the signal distribution.

Similarly to $B^+ \rightarrow K^+ \pi^+ \pi^- \eta$ decays, which contaminate the post-selection data significantly in the region of the signal peak, $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ decays present another source of background that is difficult to reduce. Neutral pions can mimic a photon signature, as mentioned in Section 5.1. The π^0 background events that survive the offline selection cannot be subtracted from the data sample with weights as the distributions of the signal and the background in the B candidate mass are so similar as to make impossible the construction of a useful orthogonal weight function for this background species: as the distinction between the two distributions

would have to be based on very minor differences in shape and position, the weight values for both signal and background move so far away from unity in the positive and negative ranges that the amplitude fit on the weighted data sample is compromised. The ratio of pions to photons can be reduced by placing a strict requirement on the IsPhoton classifier score [47] introduced in Section 5.1. However, to reduce it to a level where it is insignificant would also mean degrading the signal efficiency to a point of undue loss of statistical power. We therefore choose to place the strictest requirement on IsPhoton consistent with sufficiently high signal efficiency and to merge the π^0 background with the signal in the background subtraction. With the selection threshold on the IsPhoton score of 0.9 chosen for this analysis, the expected contamination from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events is $(1.31 \pm 0.21)\%$ for Run 1 and $(1.46 \pm 0.23)\%$ for Run 2, see Table 5.8. Similarly to what is done for the $B^+ \rightarrow K^+ \pi^+ \pi^- \eta$ background, for the calculation of these contamination values, it is assumed that the ratio of the background and signal branching fractions is the same as the ratio of branching fractions of the decays $B^+ \rightarrow K^{*+} \pi^0$ and $B^+ \rightarrow K^{*+} \gamma$. The remaining π^0 contamination is a potential source of bias on the result of the amplitude fit. Studies have been performed to better estimate the π^0 contamination using the IsPhoton distribution in data. These estimates are used to determine the systematic uncertainty on λ_γ associated with the $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ background, and are detailed in Section 6.8.3.

In the fit to the B mass distribution in data, the π^0 contamination is fixed to the result of the calculation based on the branching fractions and MC efficiencies. The $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ background is modelled, like the signal, with a double-sided CB, see Eq. 5.1. The tail parameters $\alpha_L, \alpha_R, n_L, n_R$ for the π^0 background as well as the width σ are determined from a fit to simulated $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events reconstructed as $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays. The mean μ of the π^0 background shape is pegged to the signal mean and shifted towards lower masses by the difference between the signal mean from a fit to simulated signal decays and the mean of the $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ MC fit ($0.7 \text{ MeV}/c^2$ in Run 1 and $8.6 \text{ MeV}/c^2$ in Run 2). The latter fit is shown in Figure 5.3. The $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ CB and the signal CB have different shape parameters, so their common distribution in the background subtraction is the sum of the two CB distributions with their relative normalisations fixed to the $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ contaminations from Table 5.8.

The combinatorial background distribution is modelled with a linear function whose two parameters are determined in the full B mass fit on data. No input is taken from simulation because, as explained in Section 5.1 for the multivariate classifier that distinguishes this background, it would be very computationally expensive. Fortunately, the combinatorial background is the only component that contaminates the spectrum at high B masses, since all other background contributions come from incorrectly and/or partially reconstructed B decays and therefore have mass distributions at or below the signal level. Therefore, the shape of the combinatorial background can be constrained in the fit to the data distribution when the fit range extends to high B mass values beyond the signal peak.

5.2. Fit to B candidate invariant mass




Figure 5.3 – $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ candidate mass distributions from simulated $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events for 2012 (left) and 2016 (right) MC. The fit to the distribution with a double-sided CB PDF is shown in blue. The pull plots at the bottom are described in Figure 5.2.

In partially reconstructed B decays, the loss of one pion shifts the B candidate mass distribution from the signal peak towards lower mass values by at least one pion mass, and analogously for two lost pions and two pion masses. Partially reconstructed backgrounds with one lost pion are combined into a single background species in the B mass fit, and partially reconstructed backgrounds with two missing pions make up another background species. The former is modelled, like the $B^+ \rightarrow K^+ \pi^+ \pi^- \eta (\rightarrow \gamma\gamma)$ background, with an Argus function, convolved with the signal shape to account for the detector resolution. The width of the resolution proxy is scaled from the signal width to account for the effect of the missing pion, which improves the resolution on the B mass. Naturally, the reconstructed value of the B mass cannot be more accurate in the partially reconstructed case than in the fully reconstructed case, but the presence of one more reconstructed track, with its associated uncertainties on its four-momentum, means that the spread of the B mass values is slightly higher in the fully reconstructed case. This effect on the resolution of removing one pion is expected to be approximately independent of the number of tracks in the B decay. For lack of a sizeable $B^0 \rightarrow K^+ \pi^- \pi^+ \pi^- \gamma$ MC sample, it is estimated from simulated $B^+ \rightarrow K^{*0} (\rightarrow K^+ \pi^-) \pi^+ \gamma$ events, which figure as the partially reconstructed background, and $B^0 \rightarrow K^{*0} (\rightarrow K^+ \pi^-) \gamma$ events, which figure as the signal. The MC samples for both decays are reconstructed as $B^0 \rightarrow K^+ \pi^- \gamma$. The offline selection for this analysis is used as far as it applies to this event type; notably, the HLT2 trigger line Hlt2RadiativeIncHHGamma (two hadrons and a photon) is chosen instead of the Hlt2RadiativeIncHHHGamma line (three hadrons and a photon) in Run 2 (the Run 1 trigger lines are unchanged, see Table 5.1), the vetoes and fiducial requirements on the invariant mass combinations are not applied, and the BDT classifier is not evaluated. Also, in the absence of a pion candidate with the same charge as the kaon, any and all selection requirements related specifically to this candidate are moot. In the resulting $B^0 \rightarrow K^{*0} \gamma$ sample, the distribution of the B^0 candidate mass is fitted with a double-tailed CB, while in the $B^+ \rightarrow K^{*0} \pi^+ \gamma$ sample reconstructed as $B^0 \rightarrow K^+ \pi^- \gamma$, the difference between the reconstructed B^0 candidate invariant mass and the true invariant mass of the $K^+ \pi^- \gamma$ system is fitted, also with a double-tailed CB. The ratio of the widths of the two




Figure 5.4 – Fits to constrain the shape of the partially reconstructed background with one missing pion: left is the difference between the reconstructed $B^0 \rightarrow K^+ \pi^- \gamma$ candidate mass and the true invariant mass of the $K^+ \pi^- \gamma$ system in $B^+ \rightarrow K^{*0} \pi^+ \gamma$ decays, in 2016 MC. The double-tailed CB resulting from the fit is superimposed. This shape is used as the resolution function in the fit to the $B^0 \rightarrow K^+ \pi^- \gamma$ candidate mass in the same sample shown on the right, superposed with the Argus PDF convolved with the CB resolution. The pull plots at the bottom are described in Figure 5.2.

fitted CB distributions is found to be $r = 0.92 \pm 0.06$ in 2012 MC and $r = 0.95 \pm 0.03$ in 2016 MC. These ratios are used to scale the signal width in the resolution function that is convolved with the Argus PDF in the shape of the partially reconstructed backgrounds with one missing pion. In the fit to the B mass of the $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ candidates in data, the end-point of the Argus is pegged to the signal mean minus the π^0 mass. The curvature and power are fixed to the result of a fit to the $B^+ \rightarrow K^{*0} \pi^+ \gamma$ -reconstructed-as- $B^0 \rightarrow K^+ \pi^- \gamma$ sample, where the fit function is an Argus with the end-point fixed to the B^+ mass minus the π^0 mass, convolved with the resolution function determined from the difference between the B^0 candidate mass and the true invariant mass $K^+ \pi^- \gamma$ system mentioned above. That fit, as well as the one that determines the Argus PDF curvature and power, is shown in Figure 5.4 for the 2016 MC sample.

The partially reconstructed backgrounds with two or more missing pions are modelled with a *RooPhysBkg* PDF [57], another generic shape for partially reconstructed backgrounds, which is defined as

$$f_{\text{part}}(x; m_0, \sigma, c) = \mathcal{N} \int_{-\infty}^{+\infty} x' \left(1 - \frac{x'^2}{m_0^2} \right) \Theta(m_0 - x') e^{-c_{\text{part}} x' - \frac{1}{2} \left(\frac{x-x'}{\sigma} \right)^2} dx'. \quad (5.3)$$

This function is similar to a simplified Argus PDF (see Eq. 5.2), multiplied with the Heaviside step function Θ , and convolved with a Gaussian resolution function. The end-point m_0 is pegged to the signal mean minus two times the neutral pion mass, the width σ of the Gaussian is pegged to the signal width, and the curvature c_{part} is left free.

The distribution of the B candidate mass in data is fitted, for each data-taking year, with a composite PDF comprised of the B mass shapes for the signal, the η background, the π^0 background, the combinatorial background, the partially reconstructed background with one

missing pion, and the partially reconstructed background with two missing pions. The yields (number of events) of the η and π^0 backgrounds are not free in the fit as they depend on the signal yield as described above. The free parameters of the fit are

- the signal yield N_{sig} ;
- the yield of the partially reconstructed background with one missing pion N_{misspi} ;
- the yield of the partially reconstructed background with two missing pions N_{part} ;
- the combinatorial background yield N_{comb} ;
- the signal mean μ (mu);
- the signal width σ (sigma);
- the curvature c_{part} of the partially reconstructed background with two missing pions;
- the slope a of the combinatorial background.

The data with the fitted distributions are shown in Figure 5.5 with the fit results. The latter are also tabulated in Table 5.9.

The total number of signal events is 35236 ± 198 . Of these, 34855 are in the B mass range $5000 - 6500 \text{ MeV}/c^2$ that is used for the amplitude fit.

Table 5.9 – Fit parameters obtained from the mass fit of the $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ candidates selected in the 2011, 2012, 2015, 2016, 2017, and 2018 data sets.

Fit parameter	2011	2012	2015
N_{sig}	1387 ± 37	4061 ± 68	1228 ± 36
N_{misspi}	854 ± 62	2127 ± 98	638 ± 47
N_{part}	659 ± 49	1963 ± 82	529 ± 40
N_{comb}	16.9 ± 5.5	15.2 ± 6.4	10.0 ± 4.1
$\mu [\text{MeV}/c^2]$	5263.4 ± 3.2	5280.2 ± 1.0	5263.7 ± 2.8
$\sigma [\text{MeV}/c^2]$	81.1 ± 2.5	77.3 ± 1.3	71.4 ± 2.3
$c_{\text{part}} [c^2/\text{MeV}]$	7.0 ± 3.3	2.3 ± 1.5	-3.2 ± 2.2
$a [c^2/\text{MeV}]$	0.999773 ± 0.000024	0.99836 ± 0.00024	0.99813 ± 0.00011
Fit parameter	2016	2017	2018
N_{sig}	8849 ± 100	8772 ± 99	10939 ± 110
N_{misspi}	5026 ± 147	4614 ± 143	6220 ± 157
N_{part}	5840 ± 129	5602 ± 127	6824 ± 142
N_{comb}	43 ± 11	67 ± 13	64 ± 13
$\mu [\text{MeV}/c^2]$	5264.6 ± 1.0	5273.4 ± 1.0	5272.97 ± 0.97
$\sigma [\text{MeV}/c^2]$	72.52 ± 0.87	72.00 ± 0.83	71.56 ± 0.78
$c_{\text{part}} [c^2/\text{MeV}]$	1.91 ± 0.74	1.45 ± 0.74	2.45 ± 0.63
$a [c^2/\text{MeV}]$	0.99796 ± 0.00025	0.99467 ± 0.00019	0.99048 ± 0.00024

Chapter 5. Data selection and B mass fit




Figure 5.5 – Invariant mass distribution of the $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ candidates selected in the 2011 (upper left), 2012 (upper right), 2015 (centre left), 2016 (centre right), 2017 (lower left), and 2018 (lower right) data sets. The total fit function is shown as the solid blue curve. It is the sum of the following components shown as dashed curves: $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ signal (blue), $B^+ \rightarrow K^+\pi^+\pi^-\eta$ background (green), $B^+ \rightarrow K^+\pi^+\pi^-\pi^0$ background (orange), background with one missing pion (cyan), background with at least two missing pions (magenta) and combinatorial background (red). The differences (residuals) between the numbers of bin entries and the average of the fitted blue curve over each bin in units of the uncertainty of the residuals are shown in the bottom plots, with grey bars indicating the deviation from zero and red dashed lines at ± 2 .

5.3 Background subtraction with custom orthogonal weight functions

The PDF for the amplitude fit, introduced in Section 6.1, describes only $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ signal decays. In order to correct for the presence of background events in the data samples that survive the selection requirements detailed above, the different distributions of the signal and background categories in the B candidate mass are leveraged. To this end, a weight function that depends on the B candidate mass is constructed from the result of the B mass fit on data (see Section 5.2) and used to subtract the background from the data in the amplitude fit: each event is assigned a weight so that the various backgrounds cancel out and the resulting weighted sample is effectively signal-only.

When the distributions of the signal and the backgrounds are known in a discriminating variable, one can use these distributions to disentangle the different species (the species being the signal and backgrounds) and obtain the distribution of another variable, which must be independent from the discriminating variable, for each species separately. The sPlot technique [58] is commonly used for this purpose; however, in this analysis, a more general technique known as Custom Orthogonal Weight (COW) functions [59] is used. Unlike the sPlot technique, which relies on a maximum likelihood fit, no fit is required for COWs when the distributions in the discriminating variable and the yields of the different species are known.

The B candidate invariant mass is used as the discriminating variable, and signal weights are applied to the data events to obtain background-subtracted, i.e. effectively signal-only, samples. The distributions in the B mass are determined from simulation and data as described in Section 5.2.

From the distributions f_i and yields N_i of the species $i = 1, \dots, N_{\text{species}}$, the weight for a species k is

$$w_k(m) = \frac{\sum_{i=1}^{N_{\text{species}}} C_{ki} f_i(m)}{\sum_{i=1}^{N_{\text{species}}} N_i f_i(m)}, \quad (5.4)$$

where m is the B candidate mass and the matrix C_{ki} is defined as

$$(C^{-1})_{kl} = \int dm \frac{f_k(m) f_l(m)}{\sum_{i=1}^{N_{\text{species}}} N_i f_i(m)}. \quad (5.5)$$

The weights are scaled with a common normalisation constant that accounts for the statistical power of the weighted data sample with respect to a sample without weights. The normalisation constant is given by

$$\mathcal{S} = \frac{\sum_{i=1}^{N_{\text{data}}} w_{\text{sig}}(m_i)}{\sum_{i=1}^{N_{\text{data}}} (w_{\text{sig}}(m_i))^2}, \quad (5.6)$$

so that the signal weight s_i for event i is given by $\mathcal{S} w_{\text{sig}}(m_i)$.

The characterisation of the signal and background shapes in Section 5.2 uses a wide B mass range between $4300 \text{ MeV}/c^2$ and $6500 \text{ MeV}/c^2$. Once the shapes are defined, however, the background-dominated low- B -mass region can be excluded. Therefore, the data is constrained to the B mass region between $5000 \text{ MeV}/c^2$ and $6500 \text{ MeV}/c^2$. This choice of range allows to combine into one species the partially reconstructed backgrounds with one missing pion, two missing pions, and one missing photon from the $\eta \rightarrow \gamma\gamma$ decay. With the combination of the $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ and $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ contributions into one species, this means that there are three species to consider in the weights calculation ($N_{\text{species}} = 3$):

- $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ and $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$;
- partially reconstructed backgrounds;
- combinatorial background.

For simplicity, the species that contains the signal and the π^0 background will be referred to as the signal species, or simply as “signal” for the rest of this section. It is to be able to subtract the combinatorial background that the end of the B mass range is chosen far from the signal peak; this way, there is a region dominated by combinatorial background that is assigned negative signal weights to counteract the combinatorial background events in the signal-dominated region. The yields of the partially reconstructed backgrounds are high enough that efficient background subtraction is possible with the lower limit of the background subtraction range relatively closer to the signal peak than the upper limit.

The background-subtracted samples reproduce the signal distributions of variables other than the discriminating B mass if and only if those other variables are independent of the B mass. For the purposes of the amplitude fit, the background subtraction should be valid for the five variables that describe the final state of the B decay, see Section 6.2. The correlations of these five variables with the B mass are examined on the “norm” signal MC. While a lack of correlation is a necessary rather than a sufficient condition for independence, the independence hypothesis can be strengthened by a visual inspection of the two-dimensional distributions of the discriminating variable and each of the five final-state phase-space variables. These are shown in Figure 5.6 in the B mass-constrained (see Section 6.3) case. The correlation coefficients are all below 2%, the highest value of 1.8% being reached for $m_{K^+ \pi^- \pi^+}$, and are hence considered negligible, and the two-dimensional distributions in Figure 5.6 show no indication of any kind of non-linear dependence that would escape detection by the correlation test. This validates the background subtraction using COWs with the B candidate mass as the discriminating variable.

5.3. Background subtraction with custom orthogonal weight functions




Figure 5.6 – Correlations between the B candidate mass and the five variables $m_{K^+\pi^-\pi^+}$, $m_{K^+\pi^-}$, $m_{\pi^-\pi^+}$, $\cos\theta$, and χ_{hel} after B -mass constraint in the “norm” MC. The sub-figures on the right-hand side show zooms in the most densely populated regions.

Chapter 6

Amplitude analysis

The data collected by the LHCb experiment contains information on the physical nature of the process $B^\pm \rightarrow K^\pm \pi^\mp \pi^\pm \gamma$. To extract (some of) this information, one first needs a qualitative framework in which the process can be described. This framework can then be compared to the data to obtain a quantitative description of the process within the framework. The qualitative framework used in this work to describe the decay $B^\pm \rightarrow K^\pm \pi^\mp \pi^\pm \gamma$ has been introduced in Section 2.3. The quantitative step takes the form of an amplitude analysis, or amplitude “fit” to obtain

- the decay paths that contribute to the transition from the initial state B to the final state $K\pi\pi\gamma$, *i.e.* the set of amplitudes referred to in the following as the model content;
- the relative magnitudes and phases of the amplitudes that make up the model, and;
- the photon polarisation parameter λ_γ .

Technically, a complete model of the decay is given by a set $\Omega = \{A_{k,R/L}, a_k, \phi_k, \lambda_\gamma\}$ of amplitudes, magnitudes and phases, and λ_γ ; however, we also use the word “model” to refer to the set $\{A_{k,R/L}\}$ only. Fitting this model to the data then leads to estimates of the parameters $\{a_k, \phi_k, \lambda_\gamma\}$.

The amplitude fit takes the data in the form of a list of signal decay candidates with 16 variables \mathbf{x} each (which are the four-momenta of the four final-state particles) and returns the best estimate of the magnitudes and phases and of the polarisation fraction polFrac , which is related to λ_γ as

$$\text{polFrac} = \frac{1 + \lambda_\gamma}{2}. \quad (6.1)$$

In fits on data, the polFrac parameter is blinded through the addition of an unknown offset that is randomly chosen from the range $[-0.2, 0.2]$. The offset is fixed for the whole analysis so that it cancels in the calculation of relative differences between blinded polFrac values.

Let $\mathcal{F}(\mathbf{x}|\Omega)$ be a function that describes the probability distribution of the measurable quantities \mathbf{x} given a set of parameters Ω . We wish to invert this relation, that is to say: given a set of data $\{\mathbf{x}_i\}$, we wish to find the “best” parameter set Ω to describe the data. With the likelihood function

$$\mathcal{L}(\Omega) = \prod_{i=1}^{N_{\text{data}}} \mathcal{F}(\mathbf{x}_i|\Omega), \quad (6.2)$$

the principle of maximum likelihood states that the best estimate of Ω is the one that maximises \mathcal{L} [60].

In practice, instead of maximising the likelihood, one can minimise the negative log-likelihood NLL,

$$\text{NLL}(\Omega) = -2 \ln \mathcal{L}(\Omega), \quad (6.3)$$

since the natural logarithm is a strictly monotonous function. The factor 2 is added so that differences in the NLL can be conveniently identified with the χ^2 test statistic in the case of normally distributed data. The background-subtraction weights s_i (see Section 5.3) are included as weights in the sum that the logarithm makes of the product in Eq. 6.2:

$$\text{NLL}(\Omega) = -2 \sum_{i=1}^{N_{\text{data}}} s_i \ln \mathcal{F}(\mathbf{x}_i|\Omega). \quad (6.4)$$

It is worth mentioning already, before describing the PDF \mathcal{F} in the next section, that some attributes of the amplitudes, such as masses and widths of the resonances for example, can also be determined from a minimisation of the NLL, and can thus be seen as a part of the set Ω for certain fits.

6.1 Fit function

The PDF that describes the signal decay in the fit function of the amplitude analysis is based on the expression of Eq. 2.14 for the differential decay rate derived in Section 2.3. Note that Eq. 2.14 describes B^- decays, whereas the amplitude fit PDF describes B^+ decays. In B^+ decays, the photon is predominantly right-handed in the SM. The decay amplitude for a right-handed photon is associated, in B^+ decays, with the Wilson coefficient \mathcal{C}_7 , while the amplitude for a left-handed photon is associated with \mathcal{C}'_7 . This means that for the differential decay rate for $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays, the indices R and L in Eq. 2.14 have to be swapped, while the photon polarisation parameter λ_γ defined in Eq. 2.13 remains the same. The signal function \mathcal{P}_s is proportional to the differential decay rate:

$$\mathcal{P}_s(\mathbf{x}|\Omega) = \frac{1 + \lambda_\gamma}{2} \left| \sum_k a_k e^{i\phi_k} A_{k,R}(\mathbf{x}) \right|^2 + \frac{1 - \lambda_\gamma}{2} \left| \sum_k a_k e^{i\phi_k} A_{k,L}(\mathbf{x}) \right|^2. \quad (6.5)$$

It is a function of the four-momenta \mathbf{x} of the four final-state particles. The signal function \mathcal{P}_s , the four-body phase-space density $\Phi_4(\mathbf{x})$, and the function $\xi(\mathbf{x})$, which encodes the LHCb acceptance as well as the reconstruction and selection efficiencies, make up the PDF \mathcal{F} that describes the signal data,

$$\mathcal{F}(\mathbf{x}|\Omega) = \frac{1}{\mathcal{N}_\Omega} \xi(\mathbf{x}) \mathcal{P}_s(\mathbf{x}|\Omega) \Phi_4(\mathbf{x}), \quad (6.6)$$

where the normalisation factor \mathcal{N}_Ω is given by

$$\mathcal{N}_\Omega = \int \xi(\mathbf{x}) \mathcal{P}_s(\mathbf{x}|\Omega) \Phi_4(\mathbf{x}) d\mathbf{x}. \quad (6.7)$$

The functions ξ and Φ_4 are intrinsically taken into account in the evaluation of the PDF on (simulated) data: any event generated in simulation or reconstructed from data must have four-momentum variables \mathbf{x} that are kinematically allowed, by the construction of the generator or the constraints of the reconstruction as signal. The efficiency $\xi(\mathbf{x})$ takes the value 0 for a signal decay that lies outside of the acceptance or that does not pass the selection requirements and the value 1 for a signal decay within the acceptance that passes the selection. It can also take values between 0 and 1 for a MC decay with PID efficiency weights, in which case the value of the signal function is multiplied by the PID efficiency weight.

The normalisation factor is computed numerically from simulated “norm” data as

$$\mathcal{N}_\Omega \approx \int \xi(\mathbf{x}) \mathcal{P}_{\text{norm}}(\mathbf{x}) \Phi_4(\mathbf{x}) d\mathbf{x} \times \frac{1}{N_{\text{sel}}} \sum_{j=1}^{N_{\text{sel}}} w_j \frac{\mathcal{P}_s(\mathbf{x}_j|\Omega)}{\mathcal{P}_{\text{norm}}(\mathbf{x}_j)}, \quad (6.8)$$

where $\mathcal{P}_{\text{norm}}$ is the signal function used in the generation of the “norm” MC samples, with the parameters from Table 6.3, and w_j are the MC weights, set to 1 if no weighting is applied. For the N_{sel} “norm” MC events that pass the selection, the signal function \mathcal{P}_s is evaluated, and the

Table 6.1 – Spin factors for the weak decay $B^- \rightarrow K_{\text{res},i}^- \gamma_{\text{R/L}}$. Where two signs are given, the upper sign belongs to the R and the lower sign to the L spin factor. The four-momentum $P = p_B + p_{K_{\text{res},i}}$ is the sum of the B and $K_{\text{res},i}$ four-momenta, while $q = p_B - p_{K_{\text{res},i}}$ is their difference, equal to the photon four-momentum. The Levi-Civita tensor $\epsilon_{\mu\nu\rho\sigma}$ is $+1$ (-1) for even (odd) permutations of (0123) , and 0 otherwise. The photon polarisation vector is ϵ . The $K_{\text{res},i}$ polarisation vector or tensor is ϵ' . The asterisk indicates complex conjugation.

Resonance type	J^P	$S_{i,\text{R/L}}$
Axial vector	1^+	$\mp i\epsilon_{\mu\nu\rho\sigma}\epsilon^{*\mu}\epsilon'^{*v}P^\rho q^\sigma + ((\epsilon^* \cdot \epsilon'^*) (P \cdot q) - (\epsilon^* \cdot P) (\epsilon'^* \cdot q))$
Vector	1^-	$-i\epsilon_{\mu\nu\rho\sigma}\epsilon^{*\mu}\epsilon'^{*v}P^\rho q^\sigma \pm ((\epsilon^* \cdot \epsilon'^*) (P \cdot q) - (\epsilon^* \cdot P) (\epsilon'^* \cdot q))$
Tensor	2^+	$-i\epsilon_{\mu\nu\rho\sigma}\epsilon^{*\mu}\epsilon'^{*v}\lambda P_\lambda P^\rho q^\sigma \pm \epsilon^{*\mu}P^\sigma(\epsilon'^*_{\mu\sigma}(P \cdot q) - (P_\mu \epsilon'^*_{\sigma\nu} q^\nu))$
Pseudo-tensor	2^-	$\mp i\epsilon_{\mu\nu\rho\sigma}\epsilon^{*\mu}\epsilon'^{*v}\lambda P_\lambda P^\rho q^\sigma + \epsilon^{*\mu}P^\sigma(\epsilon'^*_{\mu\sigma}(P \cdot q) - (P_\mu \epsilon'^*_{\sigma\nu} q^\nu))$

sum of the ratios of the two signal functions is computed. The integral in Eq. 6.8 and the factor $\frac{1}{N_{\text{Sel}}}$ do not change with different fit parameters Ω , and these constant factors (additive in the NLL) can therefore be omitted in the minimisation.

The propagators $\mathcal{T}_i^k(\mathbf{x})$ and $\mathcal{T}_j^k(\mathbf{x})$ of Eq. 2.15 are generally described by relativistic Breit-Wigner distributions [61]. The broad $\rho^0 \rightarrow \pi^+ \pi^-$ resonance and the narrow $\omega \rightarrow \pi^+ \pi^-$ resonance overlap and are described with a common line-shape that is adapted from the Gounaris-Sakurai line-shape [62]. Expressions for all propagators can be found in Appendix A.2.

Angular momentum conservation in the weak decay of the pseudo-scalar B meson into a photon and a kaonic resonance requires that the resonance have non-zero integer spin. No significant contributions from resonances with spin higher than 2 are expected in the invariant mass ranges considered in this analysis, so only vector ($J^P = 1^-$), axial vector ($J^P = 1^+$), tensor ($J^P = 2^+$), and pseudo-tensor ($J^P = 2^-$) resonances are considered. The decays of these resonances into the three pseudo-scalar hadrons of the final state are mediated by the strong interaction.

The weak decay spin factors $S_{i,\text{R/L}}$ from Eq. 2.15 were calculated using the helicity formalism in Ref. [20] for vector, axial vector, and tensor resonances. The expression for pseudo-tensor resonances is inferred by analogy to the tensor expression, and by requiring that the R and L amplitudes differ by a factor $P_{\text{res}}(-1)^{J_{\text{res}}-1}$. The weak decay spin factors are given in Table 6.1. They are defined in the B rest frame. A boost of the $K\pi\pi\gamma$ system in a direction other than the axis defined by the K_{res} and γ momenta in the B rest frame distorts the angular distributions. That is why the four-momenta in the data and integration samples are boosted from the LHCb lab frame to the B rest frame for the amplitude fit.

In the isobar formalism, each decay channel contains two strong decays: that of the initial resonance $K_{\text{res},i}$ into an intermediate resonance R_j and one of the final-state pseudo-scalars, and that of R_j into the rest of the final state. A strong two-body decay $R \rightarrow AB$ conserves angular momentum and parity: with the spin (parity) J_X (P_X) of the particle $X = R, A, B$, the total spin S_{AB} of the daughters can take values between $|J_A - J_B|$ and $J_A + J_B$. Conservation of

Table 6.2 – Intermediate states $K_{\text{res},i}$ and R_j taken into consideration in the model building, grouped by their possible decay products and spin-parity properties.

	$J^P = 0^+$	$J^P = 1^-$	$J^P = 1^+$	$J^P = 2^+$	$J^P = 2^-$
$K_{\text{res},i} \rightarrow R_j \pi$		$K^*(1410)$	$K_1(1270)$		
		$K^*(1680)$	$K_1(1400)$	$K_2^*(1430)$	$K_2(1580)$
$K_{\text{res},i} \rightarrow R_j K$	non-resonant	$K_1(1650)$	$K_1(1650)$		$K_2(1770)$
$R_j \rightarrow K\pi$	$K_0^*(1430)$	$K^*(892)$		$K_2^*(1430)$	
$R_j \rightarrow \pi\pi$	non-resonant	$\rho(770)$ $\omega(782)$		$f_2(1270)$	

angular momentum requires that J_R lie between $|L_{AB} - S_{AB}|$ and $L_{AB} + S_{AB}$, where L_{AB} is the relative angular momentum of the daughters. Another constraint on the spin-orbit structure of the decay follows from parity conservation: $P_R = P_A P_B (-1)^{L_{AB}}$. Note that in this analysis, all final-state hadrons are pseudo-scalar particles, so without loss of generality, we can take $J_B = 0$ and $P_B = -1$.

The centrifugal barrier factors limit the relevant angular momenta, so we only consider S-, P-, and D-wave decays (corresponding to relative angular momenta L_{AB} of 0, 1, and 2, respectively). The set of resonances and non-resonant states considered for the intermediate states $K_{\text{res},i}$ and R_j are given in Table 6.2.

There are no established axial vector nor pseudo-tensor candidates for the two-body resonance R_j , so with constraints from angular momentum and parity conservation, we consider ten spin-orbit configurations for the $K_{\text{res},i}$ decays. The particles involved are scalars S ($J^P = 0^+$), pseudo-scalars P ($J^P = 0^-$), vectors V ($J^P = 1^-$), axial vectors A ($J^P = 1^+$), tensors T ($J^P = 2^+$), and pseudo-tensors Pt ($J^P = 2^-$). The spin-orbit configurations for axial vector $K_{\text{res},i}$ are $A \xrightarrow{P} SP$, $A \xrightarrow{S} VP$, $A \xrightarrow{D} VP$, and $A \xrightarrow{P} TP$. For vector $K_{\text{res},i}$, they are $V \xrightarrow{P} VP$ and $V \xrightarrow{D} TP$. The $K_2^*(1430)$ is the only tensor considered among the resonances that decay to $K\pi\pi$. $T \xrightarrow{P} TP$ decays are allowed, but the only possible tensor for R_j in this case is the $f_2(1270)$, and conservation of energy heavily suppresses the decay $K_2^*(1430) \rightarrow f_2(1270)K$. Therefore, only $T \xrightarrow{D} VP$ decays are considered. The spin-orbit configurations considered for pseudo-tensor $K_{\text{res},i}$ are $Pt \xrightarrow{P} VP$, $Pt \xrightarrow{S} TP$, and $Pt \xrightarrow{D} SP$. The $Pt \rightarrow TP$ decay is also allowed in the D-wave configuration, but only the S-wave is considered in the model building. This is because the tensor two-body resonances $f_2(1270)$ and $K_2^*(1430)$ peak beyond the limits of the $\pi^+\pi^-$ and $K^+\pi^-$ invariant mass ranges, respectively, that are used in this analysis, although their large widths mean that they can have significant tails in the mass windows considered here. In order to simplify the model building, only the lowest accessible waves are considered for the tensor two-body resonances. All two-body resonances R_j decay to two pseudo-scalars whose relative angular momentum is determined by the mother's spin, so the spin-orbit configurations considered are $S \xrightarrow{S} PP$, $V \xrightarrow{P} PP$, and $T \xrightarrow{D} PP$.

Following the covariant formalism developed in Refs. [63, 64], the spin factor for the strong decay of a resonance R into two daughters A and B is built from the polarisation tensors of R , A , and B , spin projection operators, and angular momentum (projection) operators. It is written as

$$\epsilon_R X(J_R, L_{AB}, S_{AB}) L_{(L_{AB})} P_{(S_{AB})} \epsilon_A^* \epsilon_B^*, \quad (6.9)$$

where the spin projection operator $P_{(S_{AB})}$ projects the conjugated polarisation tensors ϵ_A^* and ϵ_B^* of the daughters onto the sub-space of total spin S_{AB} . The angular momentum tensor $L_{(L_{AB})}$ is constructed from the components of the relative momentum $p_A - p_B$ of the daughters, projected onto the sub-space of angular momentum L_{AB} by the spin projection operator $P_{(L_{AB})}$. Depending on the spin configuration, the factor $X(J_R, L_{AB}, S_{AB})$ is equal to 1 (when $J_R + L_{AB} + S_{AB}$ is even) or equal to the Levi-Civita four-tensor contracted with the four-momentum $p_R = p_A + p_B$. An analogous term $X(S_{AB}, S_A, S_B)$ has been omitted as, for any strong decay in this analysis, daughter B is always a pseudo-scalar, so $S_{AB} + S_A + S_B = S_{AB} + S_A = S_A + S_A$ is always even. Expressions for the polarisation tensors and the projection operators are given in Appendix A.3. The full strong decay spin factor $S_{ij,R/L}^k$ is obtained by multiplying the spin factors for the $K_{\text{res},i}$ and R_j decays.

The isobar formalism is applied even for decay chains that do not proceed, or only partly proceed, through hadronic resonances. Non-resonant placeholders are used in some decay amplitudes in the place of the resonance $K_{\text{res},i}$, or in the place of a two-pion resonance that decays to $\pi^+ \pi^-$. Rather than a phase-space $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decay, the decay $B^+ \rightarrow K(\text{NR}, 1^+)^+ (\rightarrow K^*(892)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$ models the decay of the B^+ into a $K^*(892)^0$, a π^+ , and a photon, where the $K^*(892)^0 \pi^+$ system has total angular momentum $J = 1$ and is in a positive parity eigenstate ($P = +1$). Decays of the B^+ into a photon and a $K^*(892)^0 \pi^+$ system with $J^P = 1^-$ (designated $K(\text{NR}, 1^-)$) are also considered, as are $\rho(770)^0 K^+$ systems with vector or axial vector properties. Analogously, the decay $K_1(1270)^+ \rightarrow a(\text{NR}, 0^+) (\rightarrow \pi^+ \pi^-) K^+$ has the $K_1(1270)^+$ decay into a K^+ and two pions that are in a scalar configuration. Instead of a Breit-Wigner propagator, a propagator that only contains the normalised Blatt-Weisskopf barrier factor $B_L(q, 0)$ calculated from the breakup momentum q , see Table A.1 in Appendix A.2, is assigned to such non-resonant particle pairs in the PDF. For B^+ decays into a non-resonant particle pair and a photon, the weak decay spin factor from Table 6.1 corresponding to the J^P of the non-resonant particle pair is used. The strong decay spin factors for non-resonant particle pairs are derived assuming parity conservation, so they are the same as the strong decay spin factors for the hadronic resonances introduced above.

6.2 Fit fractions, visualisation and goodness of fit

Fit fractions It is convenient to compare different amplitude models by comparing the fit fractions of the respective amplitudes instead of the magnitudes and phases. The fit fractions can be defined for single or multiple amplitudes that are a sub-set of the full model. For a sub-set $\{\mathcal{A}_{\ell,R/L}\}$, the fit fraction is defined as

$$\frac{1}{\int \xi(\mathbf{x}) \mathcal{D}_s(\mathbf{x}) \Phi_4(\mathbf{x}) d\mathbf{x}} \int \xi(\mathbf{x}) \left(\frac{1+\lambda_\gamma}{2} \left| \sum_{\ell} \mathcal{A}_{\ell,R} \right|^2 + \frac{1-\lambda_\gamma}{2} \left| \sum_{\ell} \mathcal{A}_{\ell,L} \right|^2 \right) \Phi_4(\mathbf{x}) d\mathbf{x}, \quad (6.10)$$

where $\mathcal{A}_{\ell,R/L} = a_\ell e^{i\phi_\ell} A_{\ell,R/L}$ and the integral in the denominator concerns the full model. In the case of different parametrisations of decay channels, which may have different normalisation conventions especially with regards to the propagators, the fit fractions give a better, model-independent, impression of the relative importance of different decay channels within the model than the relative magnitudes a_k . For sub-models that contain more than one amplitude, the fit fraction takes interferences between amplitudes into account. It includes the terms $\Re(\mathcal{A}_{m,R/L}^* \mathcal{A}_{n,R/L})$ if both amplitudes $\mathcal{A}_{m,R/L}$ and $\mathcal{A}_{n,R/L}$ are part of the sub-model. It will, however, not include interferences between $\mathcal{A}_{m,R/L}$ and $\mathcal{A}_{k,R/L}$ if the amplitude $\mathcal{A}_{k,R/L}$ is not in the sub-model. This can lead to fit fractions that do not add up to 100%, even if the set of sub-models is exclusive (no overlap between the sets) and complete (the union of the sub-sets makes up the full model). The integrals in the fit fraction definition are computed analogously to the PDF normalisation in Eq. 6.8.

Model visualisation in five independent variables The $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decay is passed to the fit PDF in the form of the four-momenta of the final-state particles. These 16 observables are not independent, though: all final-state particles are on-shell (4 constraints), the four-momentum of the initial state is conserved (4 constraints), and the orientation of the reference frame is arbitrary (3 constraints), so the final state has five independent observables. We choose to describe it with three masses and two angles (see Figure 6.1):

- $m_{K\pi\pi}$: the invariant mass of the $K^+ \pi^- \pi^+$ system;
- $m_{K\pi}$: the invariant mass of the $K^+ \pi^-$ system;
- $m_{\pi\pi}$: the invariant mass of the $\pi^+ \pi^-$ system;
- $\cos\theta$: the cosine of the angle between the normal of the hadronic decay plane that points in the direction of the cross product of the π^+ and π^- momenta and the opposite of the photon momentum in the $K\pi\pi$ rest frame;
- χ : the angle between the sum of the momenta of the two pions and the projection of the opposite of the photon momentum in the hadronic decay plane in the $K\pi\pi$ rest frame.

These five variables are calculated on the data and MC samples. With the background subtraction weights and the PID efficiency weights applied to data and MC respectively, the distributions in data can be compared with the distributions given by a certain amplitude




Figure 6.1 – Sketch of the $K^+\pi^-\pi^+\gamma$ final state in the rest frame of the three hadrons ($\vec{p}_{K^+} + \vec{p}_{\pi^+} + \vec{p}_{\pi^-} = 0$). The vector \hat{n} perpendicular to the hadronic decay plane is the normalised cross product of the two pion momenta $\vec{p}_{\pi^+} \times \vec{p}_{\pi^-}$. The angle between \hat{n} and the opposite of the photon momentum $-\vec{p}_\gamma$ is given by θ , and χ is the angle between the sum of the pion momenta $\vec{p}_{\pi^+} + \vec{p}_{\pi^-}$ and the projection of $-\vec{p}_\gamma$ onto the hadronic decay plane.

model. To this end, the MC samples have another set of weights applied (multiplicatively) to transform them from the “norm” model to the model under consideration. For each event, this model weight is calculated as the ratio of the model PDF and the generation (“norm”) PDF. Sub-sets of the decay amplitudes can also be shown in the histograms by calculating a sub-model weight where the sub-model contains only some sub-set of the amplitudes in the full model. The MC histograms for the full model are normalised to match the data, and any sub-model to match the data times the fit fraction for the sub-model. Since the fit fractions do not necessarily sum up to one, the sum of the sub-model histograms does not necessarily sum up to the histogram for the full model. For each of the five visualisation variables, the residuals, defined as the distance between the data points and the full model histogram, in units of the error on the data bins, can be computed and visualised. These “pull plots” are shown below the histograms that show the data and the model distributions. The error bars on the pull plots are combinations of the data and reweighted MC histogram errors. For visibility, only pulls between -6 and 6 are shown; note that this may mean that there are some bins whose pulls are not plotted.

Figure of merit: χ^2/ndf The figure of merit chosen for the correspondence between the model and the data is a χ^2 divided by its number of degrees of freedom (ndf) in the fit. The χ^2/ndf is computed in bins of the five-dimensional phase-space of the decay. The bins are rectangular in terms of five independent squared mass combinations of the final-state particles, namely $m_{\pi^+\pi^-\gamma}^2$, $m_{K^+\pi^-}^2$, $m_{K^+\pi^+}^2$, $m_{\pi^+\gamma}^2$, and $m_{K^+\pi^+\pi^-}^2$. The bins are determined using an adaptive binning scheme for each data-taking year separately. The binning algorithm [65] recursively splits the sample into sub-samples of equal size along one of the binning variables. The splitting stops when a further division would lead to bins containing fewer than the minimum number of events, which is set to 15 in this analysis. Note that the splitting points and the number of events are determined without taking event weights into account. For each bin, the variable that determines the next split is that in which the bin population varies

the most. The variation is quantified by the spread of nearest-neighbour distances of the events [66]. The binning can be visualised as a decision tree. For each bin, the contribution to the χ^2 is determined as the squared difference between the sum of the weights of the data events and the sum of the weights of the integration events, divided by the variance of that difference. The weights of the integration events have been scaled so that the sum of integration event weights over the whole sample is equal to the sum of data event weights over the whole sample. The variance is calculated as the sum of the squared data event weights and the squared scaled integration event weights. These squared ratios are summed over all the bins to give the χ^2 , which is divided by the number of degrees of freedom. They are given by the number of bins minus the number of free parameters in the fit minus one (the normalisation of the integration sample to the data sample size removes an extra degree of freedom). Studies in Ref. [65] indicate that the systematic precision of the χ^2/ndf determined this way is approximately 0.02.

6.3 Data and MC samples

The fit is performed on data samples recorded by the LHCb experiment in 2011, 2012, 2015, 2016, 2017, and 2018. The selection procedures described in Chapter 5 are applied on the data, including the fiducial cuts on the invariant masses of the hadronic systems in Table 5.6. For the amplitude fit, the B mass range is restricted to $5000 - 6500 \text{ MeV}/c^2$ and the background subtraction weights calculated in this region (see Section 5.3) are applied. As explained in Section 2.3, the data events that are reconstructed as B^- meson decays rather than B^+ meson decays have CP conjugation applied by inverting the charges and the three-momenta of all the particles involved. This way, all events can be treated as B^+ decays and the measurement considered CP -averaged. The necessary assumption of negligible CP violation is cross-checked by fitting the B^+ and the B^- events separately in Section 6.7.

Simulated samples, generated with AMPGEN [67] and with the LHCb detector response and reconstruction applied, are used for the normalisation of the fit PDF and for validation. Two models are used for generating two kinds of samples, referred to as “norm” for “normalisation” and “val” for “validation,” although they are also used for other purposes. In the generation, the decays of B^- mesons which make up about half of the events are treated the same way as the decays of B^+ mesons, so these samples need not have CP conjugation applied. The model composition (decay chain amplitudes) and parameters were chosen so that they resemble the true physical structure of the $B \rightarrow K\pi\pi\gamma$ decay, based on the Dalitz analysis [12] and on known properties of the (intermediate) decaying mesons. In Section 6.1, it is explained that the amplitude fits use the four-momenta of the final-state particles in the B rest frame. The exception to this are validation fits on the “val” and “norm” MC samples. Their generation uses rejection sampling: an event with four-momenta \mathbf{x} is retained if a random number sampled uniformly between 0 and 1 is smaller than $\mathcal{P}_{\text{gen}}(\mathbf{x})/\max(\mathcal{P}_{\text{gen}})$, and rejected otherwise. In the generation, $\mathcal{P}_{\text{gen}}(\mathbf{x})$ is evaluated on the true four-momenta \mathbf{x} in the LHCb lab frame, so in order to recover the “val” and “norm” model parameters, the fit PDF must be calculated in the same reference frame. This is another way of saying that because the generation of the “norm” and “val” MC samples uses the lab frame momenta, the generated models are different from what was intended. This is not a problem for the amplitude fit, however, because the generation PDF \mathcal{P}_{gen} for the numerical integration of the fit PDF (see Eq. 6.8) can be calculated correctly by simply using the true lab frame momenta. For this, the maximal value of the generation PDF, $\max(\mathcal{P}_{\text{gen}})$, which is determined prior to the generation as the maximal value of \mathcal{P}_{gen} over one million phase-space events, times a safety factor of 3, is also taken into account. Note that in the amplitude analysis, the lab frame momenta are only used to obtain the generation PDF of the MC samples. For all other purposes, *e.g.* for the evaluation of the model PDF in the numerator of Eq. 6.8, or the evaluation of the fit model PDF in the model visualisation (see Section 6.2), the true momenta of the MC events are boosted to the B rest frame.

Table 6.3 – The amplitudes of the $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ samples generated with AMPGEN and their relative magnitudes and phases. The B^+ radiative decay into the resonance at the head of the decay in the “Amplitude” column is implied, as are the two-body decays of the neutral resonances; for example, the first row implies $B^+ \rightarrow K_1(1270)^+ (\rightarrow K^*(892)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$. The lowest possible value of the relative angular momentum L of the K_{res} daughters is used, except where a [D-wave] indicates that $L = 2$. All the amplitudes listed here are part of the “norm” model. The right-most column indicates with a cross those amplitudes that are part of the “val” model. All samples were generated with $\lambda_\gamma = -1$.

J^P	Amplitude k	Magnitude a_k	Phase ϕ_k (rad)	“val”
1^+	$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$	1.	0.	×
	$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+ [\text{D-wave}]$	1.072	0.	
	$K_1(1270)^+ \rightarrow K^+ \rho(770)^0$	2.017	-0.910	×
	$K_1(1270)^+ \rightarrow K_0^*(1430)^0 K^+$	0.564	-1.637	
	$K_1(1270)^+ \rightarrow K^+ \omega(782)^0$	0.114	0.302	
1^-	$K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$	0.609	-0.755	×
	$K^*(1410)^+ \rightarrow K^*(892)^0 \pi^+$	0.716	0.	×
	$K^*(1680)^+ \rightarrow K^*(892)^0 \pi^+$	0.268	0.443	×
2^+	$K^*(1680)^+ \rightarrow K^+ \rho(770)^0$	0.193	1.403	
	$K_2^*(1430)^+ \rightarrow K^*(892)^0 \pi^+$	1.253	0.	×
	$K_2^*(1430)^+ \rightarrow K^+ \rho(770)^0$	2.246	1.798	×
2^-	$K_2^*(1430)^+ \rightarrow K^+ \omega(782)^0$	0.465	-2.353	
	$K_2(1580)^+ \rightarrow K^*(892)^0 \pi^+$	1.103	2.883	×
	$K_2(1580)^+ \rightarrow K^+ \rho(770)^0$	0.874	2.442	×
	$K_2(1770)^+ \rightarrow K^*(892)^0 \pi^+$	0.823	0.	×
	$K_2(1770)^+ \rightarrow K^+ \rho(770)^0$	0.191	2.527	
	$K_2(1770)^+ \rightarrow K_2^*(1430)^0 \pi^+$	0.169	-2.060	
	$K_2(1770)^+ \rightarrow K^+ f_2(1270)^0$	0.353	-0.174	
	Phase space	1.305	0.	

The “norm” model includes 19 amplitudes which are given, with their relative magnitudes and phases, in Table 6.3. One of these amplitudes does not include any resonances, but is a “direct” decay of the B meson into the four-particle final state, where every configuration of the final four-momenta that is allowed by the kinematic constraints of the decay is equally probable. This “phase-space” component was included in the “norm” model to ensure coverage of the phase-space regions away from resonance peaks, so that the normalisation of the fit PDF has high enough precision even for fit models that are very different from the resonant part of the “norm” model. The strangeness-violating decay $B^+ \rightarrow K_1(1270)^+ (\rightarrow K_0^*(1430)^0 (\rightarrow \pi^+ \pi^-) K^+) \gamma$ was included in the model accidentally instead of the decay $B^+ \rightarrow K_1(1270)^+ (\rightarrow K_0^*(1430)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$. With a fit fraction of 0.068%, the contribution of this amplitude is quite low. Since the generation PDF can be calculated, the validity of the “norm” samples is not compromised by the presence of this amplitude, even if the contribution were higher. Several million “norm” MC events (1 035 589, 3 427 474, 1 043 203,

5354017, 5150473, and 9145559 for years 2011, 2012, 2015, 2016, 2017, and 2018, respectively) were generated and simulated with LHCb detector response and reconstruction applied for the data-taking conditions of each year.

The “val” model contains 10 resonant amplitudes, indicated with crosses in Table 6.3. There is one “val” sample per LHC run; one, with 231851 events, generated and reconstructed in 2012 data-taking conditions, and one with 301035 events, generated and reconstructed in 2016 data-taking conditions.

The momentum reconstruction does not a priori take the known B^+ mass into account; instead, the mass of the mother is calculated from the properties of the daughters as determined from the reconstructed tracks (charged hadrons) or calorimeter clusters (photon). This unbiased B candidate mass is used in the fit described in Section 5.2. The momentum resolution for signal decays can be improved by taking the known B^+ mass into account. The “decay tree fit” [68] approach simultaneously reconstructs all the particles in the decay under the constraints that they come from a common vertex and that the mass of the mother is the B^+ mass. This approach especially improves the resolution of the photon momentum, which has no associated track. These “ B -mass-constrained” momenta are used in the amplitude fit. For the computation of the normalisation (see Eq. 6.8) the true (generated) momenta available in MC are used.

6.4 Iterative amplitude selection

To build the amplitude model that best describes the data, an iterative approach is used. The model used at the first iteration, called the base model, includes most of the decay chains found in the Dalitz analysis [12]. The specific make-up of the base model has been established during the development of the amplitude analysis and model building procedure, in the course of which a large number of configurations of the data and integration samples, *e.g.* concerning background subtraction in data and data-driven corrections to the MC samples, and model building algorithms were tested. The base model unites those decay amplitudes that were most frequently picked up early on in these model building tests and that lead to a good description of the data distributions. The full set of amplitudes in the base model is given in Table 6.4. Amplitudes with the decay channels $K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$ [D-wave] and $K^*(1410)^+ \rightarrow K^+ \rho(770)^0$, while not part of the Dalitz analysis result, are included in the base model.

As a preliminary step of the model building, the relative magnitudes and phases of the base model amplitudes, as well as the mean mass (m_0) values for the $K_1(1270)^+$ and the $K_2^*(1430)^+$, and the mean masses and decay widths (Γ_0) of the $\rho(770)^0$ and the $\omega(782)^0$ are obtained from a fit on data. These parameters are free throughout the model building. The $\rho(770)^0$ and the $\omega(782)^0$ masses and widths are constrained to the vicinity of their established values from Ref. [45] (referred to as PDG values) by adding Gaussian constraint terms to the minimisation. Only the central values of these Gaussian constraint terms are determined by the PDG values (given in Table 6.8). Setting the widths of the constraints to the experimental uncertainties of the PDG values makes the constraints too strict to allow the masses and widths to move in the fit. Therefore, widths larger than the PDG uncertainties are needed to allow some adjustment

Table 6.4 – Decay amplitudes included in the base model. The decay of the B^+ into the kaonic resonance at the head of the decay in the “Amplitude” column and a photon is implied, as are the two-body decays of the neutral resonances.

J^P	Amplitude k
1^+	$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$
	$K_1(1270)^+ \rightarrow K^+ \rho(770)^0$
	$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$ [D-wave]
	$K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$
	$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0$
1^-	$K^*(1410)^+ \rightarrow K^*(892)^0 \pi^+$
	$K^*(1410)^+ \rightarrow K^+ \rho(770)^0$
	$K^*(1680)^+ \rightarrow K^*(892)^0 \pi^+$
	$K^*(1680)^+ \rightarrow K^+ \rho(770)^0$
2^+	$K_2^*(1430)^+ \rightarrow K^*(892)^0 \pi^+$
	$K_2^*(1430)^+ \rightarrow K^+ \rho(770)^0$

to the data but retain the $\rho(770)^0$ and $\omega(782)^0$ identities. The widths of the four constraints of these resonances are all set to $0.4 \text{ MeV}/c^2$. The $K_1(1270)^+$ and $K_2^*(1430)^+$ masses are left free without constraint in order to improve the matching of the model to the data in the regions of the prominent peak and of the “shoulder” in the $m_{K\pi\pi}$ spectrum. Freeing the $K_2^*(1430)^+$ mass also adds some freedom to the $J^P = 2^+$ component of the amplitude model. The polarisation parameter is also free in the fit.

The fit of the base model to the data is the preparation of the model building (step zero). In step one, the best of the 33 amplitudes in Table 6.5 is chosen as an addition to the base model. These test amplitudes are temporarily added to the base model one at a time, and each of these models with 12 amplitudes (where the base model contains 11 amplitudes) is fitted to the data. The test amplitudes are sorted by how much they improve the description of the data using the χ^2/ndf metric introduced in Section 6.2. Since the uncertainty on the value of χ^2/ndf is 0.02, all test amplitudes whose fits yield a χ^2/ndf within 0.02 of the minimum over all the test amplitudes are short-listed as candidates for addition to the base model. Among these, the candidate with the most significant fraction (where significance is defined as the fit fraction divided by the fit fraction uncertainty) is permanently added to the base model, marking the end of step one. Each subsequent step of the model building uses the same selection procedure to add another decay amplitude to the model. We stop adding amplitudes to the model when the χ^2/ndf stabilises, *i.e.* when additional decay amplitudes no longer improve the χ^2/ndf value by at least 0.02. Once the χ^2/ndf has stabilised, any amplitudes with a (recalculated) fit fraction significance below 1 are removed from the model in order to reduce the model complexity. The remaining amplitudes define the nominal model resulting from the model building.

The model is then further refined by running fits to data where the masses of the kaonic resonances and the masses and widths of the $\rho(770)^0$ and the $\omega(782)^0$ are optimised in turn. In this “post-production” of the nominal model, the magnitudes and phases of the amplitudes are left free in all the fits on data. In a first step, the $\rho(770)^0$ and the $\omega(782)^0$ masses and widths are fixed to their values from the iterative model building and the K_{res} masses that were fixed to their PDG values in the iterative model building are freed, but constrained to the vicinity of their PDG values by adding Gaussian constraints. The $K_1(1270)^+$ and $K_2^*(1430)^+$ masses remain free, but are now also constrained with Gaussian penalty terms. Conversely to the case of the $\rho(770)^0$ and $\omega(782)^0$ masses and widths explained above, setting the widths of the Gaussian constraint terms of the K_{res} masses to the experimental uncertainties of the respective values in the PDG made the constraints too weak to stabilise the fit. For the K_{res} masses, a Gaussian constraint width of $4 \text{ MeV}/c^2$ was found to allow an adaptation of the masses to better match the $K\pi\pi$ system mass in data while still keeping the model description sound. After this fit on data, the K_{res} masses are fixed to the fit results, and the $\rho(770)^0$ and $\omega(782)^0$ masses and widths are freed, with Gaussian constraints to their PDG values of width $0.4 \text{ MeV}/c^2$, as in the iterative model building. In the last step of the post-production, all masses and widths are fixed and the fit on data is re-run.

Table 6.5 – Additional amplitudes tested in the model building. The B^+ radiative decay into the kaonic resonance at the head of the decay in the “Amplitude” column is implied, as are the two-body decays of the neutral resonances. The lowest possible value of the relative angular momentum L of the K_{res} daughters is used, except where a [D-wave] indicates that $L = 2$.

J^P	Amplitude k
1^+	$K_1(1270)^+ \rightarrow K^+ \rho(770)^0$ [D-wave]
	$K_1(1270)^+ \rightarrow K_0^*(1430)^0 \pi^+$
	$K_1(1270)^+ \rightarrow K^+ a(\text{NR}, 0^+)$
	$K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$ [D-wave]
	$K_1(1400)^+ \rightarrow K^+ \rho(770)^0$
	$K_1(1400)^+ \rightarrow K^+ \rho(770)^0$ [D-wave]
	$K_1(1400)^+ \rightarrow K_0^*(1430)^0 \pi^+$
	$K_1(1400)^+ \rightarrow K^+ a(\text{NR}, 0^+)$
	$K(\text{NR}, 1^+)^+ \rightarrow K^*(892)^0 \pi^+$
	$K(\text{NR}, 1^+)^+ \rightarrow K^*(892)^0 \pi^+$ [D-wave]
	$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0$ [D-wave]
	$K(\text{NR}, 1^+)^+ \rightarrow K_0^*(1430)^0 \pi^+$
	$K(\text{NR}, 1^+)^+ \rightarrow K^+ a(\text{NR}, 0^+)$
	$K(\text{NR}, 1^+)^+ \rightarrow K^+ f_2(1270)^0$
	$K(\text{NR}, 1^+)^+ \rightarrow K_2^*(1430)^0 \pi^+$
1^-	$K^*(1680)^+ \rightarrow K^+ f_2(1270)^0$
	$K^*(1680)^+ \rightarrow K_2^*(1430)^0 \pi^+$
	$K(\text{NR}, 1^-)^+ \rightarrow K^*(892)^0 \pi^+$
	$K(\text{NR}, 1^-)^+ \rightarrow K^+ \rho(770)^0$
	$K(\text{NR}, 1^-)^+ \rightarrow K^+ f_2(1270)^0$
	$K(\text{NR}, 1^-)^+ \rightarrow K_2^*(1430)^0 \pi^+$
2^-	$K_2^*(1580)^+ \rightarrow K^*(892)^0 \pi^+$
	$K_2^*(1580)^+ \rightarrow K^+ \rho(770)^0$
	$K_2^*(1580)^+ \rightarrow K_0^*(1430)^0 \pi^+$
	$K_2^*(1580)^+ \rightarrow K^+ a(\text{NR}, 0^+)$
	$K_2^*(1580)^+ \rightarrow K^+ f_2(1270)^0$
	$K_2^*(1580)^+ \rightarrow K_2^*(1430)^0 \pi^+$
	$K_2^*(1770)^+ \rightarrow K^*(892)^0 \pi^+$
	$K_2^*(1770)^+ \rightarrow K^+ \rho(770)^0$
	$K_2^*(1770)^+ \rightarrow K_0^*(1430)^0 \pi^+$
	$K_2^*(1770)^+ \rightarrow K^+ a(\text{NR}, 0^+)$
	$K_2^*(1770)^+ \rightarrow K^+ f_2(1270)^0$
	$K_2^*(1770)^+ \rightarrow K_2^*(1430)^0 \pi^+$

6.5 Validation of the amplitude fit method

The principle of the sensitivity of $B \rightarrow K\pi\pi\gamma$ decays to the amplitude model including the photon polarisation parameter λ_γ has been presented in Section 2.3. Bellée *et al.* have shown in Ref. [13] the feasibility of an amplitude fit of this decay to measure λ_γ in the absence of experimental (reconstruction) effects. The weak decay spin factors used in the study were calculated in a covariant formalism, so they differ from those used in this work which follow the helicity formalism. However, since the spin factors in the study obey the same relation between right-handed and left-handed amplitudes (a sign factor of $P_{\text{res}}(-1)^{J_{\text{res}}-1}$) discussed in Section 2.3, the conclusion on the sensitivity remains valid. The new spin factors were implemented as an optional parametrisation in the AMPGEN framework [67] for use in this analysis. The validation of AMPGEN is documented in Ref. [65]. The self-consistency of the AMPGEN framework in the context of this analysis is tested using 60 samples of 30 000 $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ events generated in AMPGEN with the “norm” model. Fitting the same model content to the samples, with the relative magnitudes, phases, and λ_γ as free parameters (37 in total), yields a set of 60 models which can be compared to the “norm” model.

Figures 6.2 and 6.3 show the distributions of the differences between the parameters and fit fractions of the fitted model and the “norm” model values, divided by the statistical uncertainties determined in the fits. These distributions, known as “pulls”, are expected to be centred at 0 if the result is unbiased, with a standard deviation of 1 if the uncertainty is estimated correctly. While the relative magnitudes and phases are not free from biases, the polFrac parameter shows neither bias nor over- or under-coverage of the statistical uncertainty.

Experimental effects from event reconstruction and selection are examined using the “val” MC samples with the full LHCb reconstruction and analysis selection (described in Section 5.1) applied. The effects of shaping of the phase-space of the fitted samples by the acceptance and efficiency of the reconstruction are assessed by fitting the “val” MC samples using the true (generated) particle four-momenta. The effects of the imperfect reconstruction of particle momenta are investigated by using the B -mass-constrained (see Section 6.3) four-momenta instead of the true four-momenta. Both “val” MC samples (2012 and 2016) are fitted simultaneously with the amplitudes from the “val” model. The fitted values of the parameters, together with their true (generation) values, are shown in Table 6.6. The fit results from true and B -mass-constrained momenta are compatible with each other. For both fit results, each fitted parameter is in the range of the true value plus or minus one or, in a few cases, two times the statistical uncertainty on the fitted parameter, except for the polFrac parameter. Here, the fit with B -mass-constrained momenta indicates a small bias that will be quantified with a systematic uncertainty on the final result, see Section 6.8.

The model building algorithm (without the post-production) is likewise checked on the “val” MC samples, using the true four-momenta of the final-state particles. The resulting model is close to the generation model, although the fitted model contains the decay chain $B^+ \rightarrow K_2^*(1430)^+ \left(\rightarrow K^*(892)^0 \left(\rightarrow K^+\pi^- \right) \pi^+ \right) \gamma$ only in the D-wave configuration, where the

generation model has only an S-wave. Notwithstanding these slight differences between the generation model and the fitted model, the result for the polFrac parameter of 0.002 ± 0.003 is compatible with the generation value of 0.




Figure 6.2 – Distribution (from 60 samples generated with “norm”) of the difference between the fit fractions of the fitted and the generated models, divided by the uncertainty on the fit fraction of the fitted model, for each of the 19 amplitudes, and for the sum of the 19 fit fractions.

Chapter 6. Amplitude analysis




Figure 6.3 – Distribution (from 60 samples generated with “norm”) of the difference between the fitted and the generated values, divided by the uncertainty on the fitted value, for each of the 37 free fit parameters.

6.5. Validation of the amplitude fit method

Table 6.6 – Model parameters and χ^2/ndf resulting from the simultaneous fits to the 2012 and the 2016 “val” samples using the generated (TRUE) and the B -mass-constrained (MACO) four-momenta, as well as the true (gen) values of the parameters of Table 6.3 used in the generation of the “val” samples.

Amplitude k		Magnitude a_k	Phase ϕ_k (rad)	Fraction (%)
$B^+ \rightarrow K_2^*(1430)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	TRUE	2.279 ± 0.067	1.827 ± 0.031	40.3 ± 0.3
	MACO	2.273 ± 0.068	1.823 ± 0.033	40.2 ± 0.3
	gen	2.246	1.798	40.2
$B^+ \rightarrow K_2^*(1430)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	1.264 ± 0.038	0.032 ± 0.032	17.4 ± 0.3
	MACO	1.256 ± 0.038	0.028 ± 0.034	17.3 ± 0.3
	gen	1.253	0.000	17.5
$B^+ \rightarrow K_1(1270)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	TRUE	2.051 ± 0.064	5.414 ± 0.029	16.7 ± 0.2
	MACO	2.045 ± 0.065	5.412 ± 0.030	16.7 ± 0.2
	gen	2.017	5.373	16.4
$B^+ \rightarrow K_2(1580)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	TRUE	0.883 ± 0.028	2.481 ± 0.035	7.3 ± 0.2
	MACO	0.884 ± 0.028	2.472 ± 0.036	7.4 ± 0.2
	gen	0.874	2.442	7.4
$B^+ \rightarrow K_2(1580)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	1.112 ± 0.038	2.906 ± 0.037	7.3 ± 0.2
	MACO	1.110 ± 0.038	2.904 ± 0.038	7.3 ± 0.2
	gen	1.103	2.883	7.4
$B^+ \rightarrow K_1(1270)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	1.000 (fixed)	0.000 (fixed)	4.3 ± 0.2
	MACO	1.000 (fixed)	0.000 (fixed)	4.3 ± 0.3
	gen	1.000	0.000	4.4
$B^+ \rightarrow K_2(1770)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	0.854 ± 0.036	0.018 ± 0.047	2.4 ± 0.1
	MACO	0.850 ± 0.037	0.009 ± 0.049	2.4 ± 0.1
	gen	0.823	0.000	2.3
$B^+ \rightarrow K_1(1400)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	0.577 ± 0.024	5.525 ± 0.049	1.9 ± 0.1
	MACO	0.583 ± 0.025	5.524 ± 0.050	1.9 ± 0.2
	gen	0.609	5.528	2.1
$B^+ \rightarrow K^*(1410)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	0.717 ± 0.036	0.023 ± 0.050	1.6 ± 0.1
	MACO	0.716 ± 0.036	0.022 ± 0.052	1.6 ± 0.1
	gen	0.716	0.000	1.6
$B^+ \rightarrow K^*(1680)^+ (\rightarrow K^*(892)^0 \pi^+) \gamma$	TRUE	0.278 ± 0.025	0.473 ± 0.098	0.35 ± 0.06
	MACO	0.279 ± 0.025	0.470 ± 0.101	0.35 ± 0.06
	gen	0.268	0.443	0.34
polFrac	TRUE	-0.000 ± 0.003		
	MACO	0.007 ± 0.003		
	gen	0.0		
χ^2/ndf	TRUE	0.98		
	MACO	1.0		

6.6 Results

The amplitudes added in each step of the model building on the experimental data and the χ^2/ndf of the respective models are collected in Table 6.7. The χ^2/ndf stabilises after step 5 with the addition of the decay amplitude $B^+ \rightarrow K_2(1770)^+ (\rightarrow K^+ \rho^0 (\rightarrow \pi^+ \pi^-)) \gamma$. In order to balance model complexity and goodness of fit, the amplitudes from steps 6 to 9 are not included in the nominal model result. Figure 6.4 shows the development of χ^2/ndf and the blinded value of the polFrac parameter, both of which stabilise after step 5.

The final “pruning” step would remove amplitudes with a fit fraction significance below 1, but no such amplitudes are found; the amplitude with the lowest significance (and the lowest fit fraction) is the one for the decay $B^+ \rightarrow K^*(1410)^+ (\rightarrow K^+ \rho^0 (\rightarrow \pi^+ \pi^-)) \gamma$ which has a fit fraction of $(0.08 \pm 0.06)\%$ and a significance of 1.3.

The post-production of the model, in which the masses and widths of the K_{res} and the light-flavour resonances are alternately optimised, leads to a model with a blinded polFrac value of 0.933 ± 0.007 and a χ^2/ndf of 1.18, both within the statistical uncertainty of the values of 0.931 ± 0.007 and 1.19 of the model before the post-production. The masses and widths that at any point in the model building are not fixed to their PDG values are shown in Table 6.8 after the iterative model building, as well as after the post-production.

The masses of both the $K_1(1270)^+$ and the $K_2^*(1430)$, free in all iterations, are each about $20 \text{ MeV}/c^2$ above their established values. The $K_1(1400)^+$ mass likewise moves about $20 \text{ MeV}/c^2$ away from its PDG value, though in the other direction, once it is freed. The other fitted K_{res} masses are compatible with their PDG values. The widths of the Gaussian constraints

Table 6.7 – Amplitudes added to the base model in the first ten steps of the iterative model building, and their χ^2/ndf values. The nominal model is developed out of the model that results from step 5.

Step	Amplitude added	χ^2/ndf	Blinded polFrac
0	None (base model)	1.37	0.899 ± 0.007
1	$B^+ \rightarrow K(\text{NR}, 1^+)^+ (\rightarrow K^+ \rho(770)^0 (\rightarrow \pi^+ \pi^-)) \gamma$ [D-wave] γ	1.35	0.912 ± 0.007
2	$B^+ \rightarrow K_1(1270)^+ (\rightarrow K_0^*(1430)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$	1.30	0.935 ± 0.007
3	$B^+ \rightarrow K(\text{NR}, 1^-)^+ (\rightarrow K^+ \rho(770)^0 (\rightarrow \pi^+ \pi^-)) \gamma$	1.28	0.924 ± 0.007
4	$B^+ \rightarrow K_1(1400)^+ (\rightarrow K^+ \rho(770)^0 (\rightarrow \pi^+ \pi^-)) \gamma$	1.22	0.914 ± 0.007
5	$B^+ \rightarrow K_2(1770)^+ (\rightarrow K^+ \rho(770)^0 (\rightarrow \pi^+ \pi^-)) \gamma$	1.19	0.931 ± 0.007
6	$B^+ \rightarrow K_1(1270)^+ (\rightarrow K^+ \rho(770)^0 (\rightarrow \pi^+ \pi^-)) \gamma$ [D-wave] γ	1.19	0.929 ± 0.007
7	$B^+ \rightarrow K(\text{NR}, 1^+)^+ (\rightarrow K^*(892)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$ [D-wave] γ	1.19	0.929 ± 0.007
8	$B^+ \rightarrow K_1(1400)^+ (\rightarrow K^*(892)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$ [D-wave] γ	1.19	0.928 ± 0.007
9	$B^+ \rightarrow K(\text{NR}, 1^+)^+ (\rightarrow K_0^*(1430)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$	1.18	0.928 ± 0.007
10	$B^+ \rightarrow K_2(1770)^+ (\rightarrow K^*(892)^0 (\rightarrow K^+ \pi^-) \pi^+) \gamma$ [D-wave] γ	1.18	0.928 ± 0.007




Figure 6.4 – Evolution of χ^2/ndf (left) and blinded polFrac (right) in the iterative part of the model building. The amplitude added at the end of each iteration (from left to right) is indicated in shorthand on the horizontal axis, where *e.g.* “1270 to RhoD” stands for the amplitude $B^+ \rightarrow K_1(1270)^+ (\rightarrow K^+ \rho^0 (\rightarrow \pi^+ \pi^-))$ [D-wave] γ . The step defining the nominal model, step 5 (note that the model labelled “base_model” is counted as step 0), is indicated in red. See also Table 6.7.

Table 6.8 – Mean masses m_0 and widths Γ_0 from the model after step 5 of the iterative part of the model building, and in the nominal amplitude model after the post-production, compared with the PDG values [45]. All values are in GeV/c^2 . The relative magnitudes and phases of the nominal amplitude model are listed in Table 6.9.

	Model after step 5	Model after post-production	PDG
$m_0(K_1(1270)^+)$	1.268 ± 0.001	1.269 ± 0.001	1.253 ± 0.007
$m_0(K_2^*(1430)^+)$	1.449 ± 0.001	1.447 ± 0.001	1.4273 ± 0.0015
$m_0(K_1(1400)^+)$	fixed to PDG	1.380 ± 0.004	1.403 ± 0.007
$m_0(K^*(1410)^+)$	fixed to PDG	1.420 ± 0.004	1.414 ± 0.015
$m_0(K^*(1680)^+)$	fixed to PDG	1.710 ± 0.004	1.718 ± 0.018
$m_0(K_2(1770)^+)$	fixed to PDG	1.769 ± 0.004	1.773 ± 0.008
$m_0(\rho(770)^0)$	0.7757 ± 0.0004	0.7757 ± 0.0004	0.77526 ± 0.00023
$\Gamma_0(\rho(770)^0)$	0.1477 ± 0.0004	0.1477 ± 0.0004	0.1474 ± 0.0008
$m_0(\omega(782)^0)$	0.7832 ± 0.0003	0.7832 ± 0.0003	0.78266 ± 0.00013
$\Gamma_0(\omega(782)^0)$	0.0107 ± 0.0003	0.0107 ± 0.0003	0.00868 ± 0.00013

of $4 \text{ MeV}/c^2$ for the K_{res} masses and $0.4 \text{ MeV}/c^2$ for the masses and widths of the $\rho(770)^0$ and $\omega(782)^0$ strongly influence the fit uncertainties on these parameters. The fit uncertainties of the masses of the $K_1(1270)^+$ and the $K_2^*(1430)^+$, which are free and unconstrained in the iterative model building, do not increase in the post-production, where they are still free, but Gaussian constraints of width $4 \text{ MeV}/c^2$ are added to the fit. The $\rho(770)^0$ and the $\omega(782)^0$ masses and widths do not change significantly in the post-production. Their deviations from the PDG values depend on the line-shape parametrisation: a model similar to the nominal one, but with an alternative description of the light flavour resonances (see Section 6.8.6) fits parameters that are compatible with the PDG values, except for the width of the $\omega(782)^0$, which, at $0.0093 \pm 0.0004 \text{ GeV}/c^2$, is nevertheless closer to the PDG value than the nominal model. Note that the fit uncertainty on this parameter is driven by (and its rounded value is equal to) the width of its Gaussian constraint. A fit of a model containing the same decay amplitudes as the nominal one, but with all masses and widths fixed to their PDG values, returns a blinded polFrac value of 0.935 ± 0.007 for a χ^2/ndf value of 1.25, compared to the value of 1.18 for the nominal model after post-production. This suggests that the description of the data is improved by freeing the masses in ways that have no strong influence on the polarisation measurement.

The relative magnitudes and phases and the fit fractions of the best amplitude model to describe the data, which is the nominal model with all the masses and widths fixed to the values obtained in the post-production step, are tabulated in Table 6.9. The corresponding projections of the data and the model on the five visualisation variables are shown in Figure 6.5.

The statistical power of the data sample is reduced through the background subtraction weights and is approximately equivalent to that of a pure sample of 28898 signal events. The sensitivity study of this method for measuring the photon polarisation, discussed in detail in Ref. [69], predicted a statistical uncertainty on λ_γ of 0.018 for a pure sample of 14 000 signal events. Scaling for sample size, this would correspond to an uncertainty of 0.013 for this measurement. Also suggested in the same discussion was a scaled prediction of this statistical uncertainty based on the uncertainty of the signal yield (before applying or accounting for background subtraction), which would lead to a statistical uncertainty of 0.012 for this measurement. The statistical uncertainty of 0.014 on λ_γ in this measurement is therefore only slightly higher than expected based on the signal yield, although the signal yield itself is lower than expected, due to the necessity of applying rather strict selection requirements. The strict selection enhances the signal purity, however, so the systematic uncertainties related to the presence of backgrounds are reduced.

The comparison of the results of this analysis for the make-up of the $K^+\pi^-\pi^+$ spectrum to other studies that include these particles in the final state is generally not straightforward. Even ignoring process dependence, differences in models and phase space (which is shaped by differences in the experimental set-up or the data selection) mean that the fit fractions of one analysis cannot be compared one-to-one with those of another analysis. The amplitude analysis of $B^+ \rightarrow \psi(2S)K^+\pi^+\pi^-$ decays [70] includes a detailed study of the $K_1(1270)^+$

substructure that translates the amplitude models from analyses of $D^0 \rightarrow K^+ K^- \pi^+ \pi^-$ [71], $\bar{D}^0 \rightarrow K^+ \pi^- \pi^+ \pi^-$ [72], and $B_s^0 \rightarrow D_s^- K^+ \pi^+ \pi^-$ [73] decays to the $B^+ \rightarrow \psi(2S) K^+ \pi^+ \pi^-$ amplitude model in order to more directly compare the fit fractions. The translated fit fractions are more in agreement than the untranslated ones. Without performing a similar study, we can still observe that the substructure of the $K_1(1270)^+$ in this analysis shows a similar ratio of contributions including $K^*(892)^0$ to contributions including $\rho(770)^0$ as the study in Ref. [70] and as the Dalitz analysis of $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decays [12] in the sense that both the fit fractions of the corresponding amplitudes are of similar magnitude, while the $K^+ \rho(770)^0$ intermediate state dominates. Contrary to the analyses mentioned in Ref. [70], including the $B^+ \rightarrow \psi(2S) K^+ \pi^+ \pi^-$ amplitude analysis, this analysis finds no significant non-resonant contribution to $K_1(1270)^+$ decays, but does find a relatively high fraction of B^+ decays to a non-resonant state in an axial vector configuration, which then decays to $K^+ \pi^- \pi^+$ via $K^+ \rho(770)^0$ intermediate states in the S- and D-wave. This analysis, the Dalitz analysis, and the amplitude analysis of $B^+ \rightarrow \psi(2S) K^+ \pi^+ \pi^-$ decays all find a contribution of the $K_1(1400)^+$ that is significantly smaller than the dominating axial vector resonance $K_1(1270)^+$. It is worth noting that both this analysis and the amplitude analysis of $B^+ \rightarrow \psi(2S) K^+ \pi^+ \pi^-$ decays find significant fit fractions for the $K_1(1400)^+ \rightarrow K^+ \rho(770)^0$ decay, where the Dalitz analysis and the PDG have clear dominance of the $K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$ channel. It is remarkable that the contributions from pseudo-tensor resonances in the Dalitz analysis is much larger than in this analysis, though a large part of this may be accounted for by the different phase-space, with the $m_{K\pi\pi}$ range of the Dalitz analysis extending by 100 MeV/c² beyond that of this analysis.

Chapter 6. Amplitude analysis

Table 6.9 – The relative magnitudes and phases of the amplitudes in the nominal decay model after the post-production as well as their fit fractions. The magnitudes and phases of the B^+ decays are relative to those of the decay $B^+ \rightarrow K_1(1270)^+ \gamma$. The magnitudes and phases of the K_{res} decays are relative to the decay of that K_{res} that has a fixed magnitude (phase) of 1 (0). The fit fractions in the rightmost column refer to the decay in the “Amplitude” column on the same line. For particles with more than one decay channel, the sum of the fit fractions is given.

polFrac (blind)	0.933 \pm 0.007		
Amplitude k	Magnitude a_k	Phase ϕ_k (rad)	Fraction (%)
$B^+ \rightarrow K_1(1270)^+ \gamma$	1.000 (fixed)	0.000 (fixed)	40.1 \pm 1.0
$K_1(1270)^+ \rightarrow K^+ \rho(770)^0$	1.469 \pm 0.050	5.311 \pm 0.038	44.8 \pm 1.6
$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	19.7 \pm 0.9
$K_1(1270)^+ \rightarrow K_0^*(1430)^0 \pi^+$	0.413 \pm 0.022	3.513 \pm 0.049	7.4 \pm 0.7
$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+ [D]$	2.481 \pm 0.117	2.460 \pm 0.050	4.2 \pm 0.3
Sum($K_1(1270)^+$)			76.1 \pm 1.2
$B^+ \rightarrow K_2^*(1430)^+ \gamma$	1.666 \pm 0.059	1.279 \pm 0.049	7.3 \pm 0.3
$K_2^*(1430)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	62.4 \pm 1.8
$K_2^*(1430)^+ \rightarrow K^+ \rho(770)^0$	0.661 \pm 0.030	0.120 \pm 0.056	24.8 \pm 1.6
Sum($K_2^*(1430)^+$)			87.2 \pm 0.3
$B^+ \rightarrow K^*(1680)^+ \gamma$	0.543 \pm 0.022	4.354 \pm 0.069	6.0 \pm 0.4
$K^*(1680)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	51.5 \pm 3.6
$K^*(1680)^+ \rightarrow K^+ \rho(770)^0$	0.612 \pm 0.055	4.010 \pm 0.122	28.0 \pm 3.1
Sum($K^*(1680)^+$)			79.5 \pm 0.7
$B^+ \rightarrow K(\text{NR}, 1^-)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	0.672 \pm 0.042	1.195 \pm 0.070	4.8 \pm 0.5
$B^+ \rightarrow K(\text{NR}, 1^+)^+ \gamma$	0.508 \pm 0.029	1.208 \pm 0.080	4.7 \pm 0.4
$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0$	1.000 (fixed)	0.000 (fixed)	91.9 \pm 1.5
$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0 [D]$	0.914 \pm 0.073	0.706 \pm 0.073	11.4 \pm 1.6
Sum($K(\text{NR}, 1^+)^+$)			103.3 \pm 0.4
$B^+ \rightarrow K_1(1400)^+ \gamma$	0.441 \pm 0.024	5.368 \pm 0.053	4.5 \pm 0.3
$K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	45.7 \pm 4.1
$K_1(1400)^+ \rightarrow K^+ \rho(770)^0$	0.742 \pm 0.070	0.823 \pm 0.090	40.2 \pm 4.0
Sum($K_1(1400)^+$)			85.9 \pm 0.4
$B^+ \rightarrow K^*(1410)^+ \gamma$	0.474 \pm 0.025	3.719 \pm 0.073	1.5 \pm 0.1
$K^*(1410)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	105.1 \pm 3.0
$K^*(1410)^+ \rightarrow K^+ \rho(770)^0$	0.226 \pm 0.085	1.365 \pm 0.298	5.2 \pm 3.9
Sum($K^*(1410)^+$)			110.3 \pm 5.4
$B^+ \rightarrow K_2(1770)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	0.809 \pm 0.040	2.204 \pm 0.071	1.2 \pm 0.1
Sum(B^+)			70.0 \pm 0.9




Figure 6.5 – Projections on the five visualisation variables (see Section 6.2) of the data (black points) and MC integration samples reweighted to the best fit amplitude model (pink) and the sub-models for the different kaonic resonances (colours indicated in the plot legends). The pulls (see Section 6.2) are shown below each histogram in the range between -6 and 6 ; this range excludes one bin each for the $m_{K\pi\pi}$ and the $m_{\pi\pi}$ plots, each of them in a region where the shapes of the distributions vary strongly between adjacent bins.

6.7 Cross-checks

Several cross-checks have been done to test the stability of the measurement, and to justify some a priori assumptions. They consist in fitting subsets of the data and/or integration samples, or alternative samples, to models with the same amplitude content as the nominal. Unless stated otherwise, the resonance masses and widths in these fits are fixed to the values from the nominal model after post-production.

The assumption that CP violation is negligible for the purposes of this analysis is cross-checked by fitting separately the data events identified as B^+ decays and those identified as B^- decays. The full MC integration sample is used in these fits for both B^+ and B^- since there is no CP violation in MC (see Section 6.3). The blinded polFrac values from the two fits are compatible: 0.934 ± 0.010 for the B^+ data subset and 0.930 ± 0.010 for the conjugated B^- data subset. This supports the assumption of negligible CP violation.

Differences in the data-taking and reconstruction over time that are not accounted for in the treatment of the samples, or not reproduced in the integration samples, could lead to a variation of the measurement results over data taken from different time periods. This is checked by fitting separately the data and integration samples from Run 1 and Run 2, as well as the data and integration samples from each individual year. The statistical power of the samples from 2011 and 2015 is too low to allow successful minimisation of the NLL, but for the other years, the blinded polFrac values are compatible. They are collected, together with the results of the other split-samples cross-checks, in Figure 6.7. The results for the Run 1 only and the Run 2 only samples have a discrepancy of 1.9 times the statistical error. However, the result for year 2012, which dominates the Run 1 sample, is compatible within one standard deviation with the Run 2 result. A fit on the samples from all years except 2011 leads to a blinded polFrac value of 0.937 ± 0.007 , 0.004 higher than the nominal value. This could indicate a bias introduced by the 2011 sample, but the effect is below the statistical uncertainty. As there is no other reason to mistrust the data recorded in 2011, the 2011 sample is retained as part of the full nominal data sample. For completeness, a fit to the data excluding the sample from 2015 yields a blinded polFrac value that is less than one permill away from the nominal value. We consider all discrepancies between the results from split samples to be due to statistical fluctuations. These cross-checks do not indicate any problem with the nominal result from the full data and integration sample sets.

As mentioned in Chapter 3, the polarity of the magnetic field of the LHCb dipole is periodically reversed to help cancel out detection asymmetries of charged particles. Any such asymmetries that are not well reproduced in simulation can mean that the “MagUp” data fits a different result than the “MagDown” data. In fact, an asymmetry between these two data sets, even if perfectly reproduced in MC, can still bias the result if it means that MC events with “MagUp” are not suitable as part of an integration sample for “MagDown” data. In this case, the fits to the two samples of data and MC split by the magnet polarity would give results that are compatible with each other, but incompatible with the result from the full data sets where the




Figure 6.7 – Blinded polFrac values for cross-checks on subsamples of the full data set: splits by magnet polarity, charge, data-taking run, and data-taking year. The result from the full data set is given in black. Note that this set of values is not statistically independent due to the overlap between samples from different splits.

polarities are mixed. The fit results from “MagUp” and “MagDown” for the blinded polFrac parameter are 0.923 ± 0.010 and 0.948 ± 0.010 , respectively, see also Figure 6.7. The two values differ by 1.8 times the statistical error. This could be due to reconstruction asymmetries that are not well accounted for in simulation, but the discrepancy is not high enough to lend much confidence to such a claim. Since the average of the two values is compatible with the nominal blinded polFrac value, these results do not indicate that the data and MC samples must be split by polarity in the fit to the full sets.

The background subtraction with custom orthogonal weight functions introduced in Section 5.3 is cross-checked by using an alternate method of reducing the background contamination in the data: a restriction of the B mass range. As the fits to the B candidate mass show (see Figure 5.5 in Section 5.2), the region of high B mass is nearly free from partially reconstructed backgrounds, while the level of combinatorial background events is low throughout. Therefore, a set of data events from that region has a low contamination of backgrounds without having to apply any background subtraction weights. The range from $5300 \text{ MeV}/c^2$ to $5600 \text{ MeV}/c^2$ is chosen for this cross-check. It comprises a little less than half of the signal peak in the B mass. The data and MC samples that are restricted to this range are therefore referred to as the “upper half peak” samples. The abovementioned fits to the full B candidate mass distribution indicate a background contamination of around 3.7% in this range, compared to over 20% in the range from $5000 \text{ MeV}/c^2$ to $6500 \text{ MeV}/c^2$ in which the background-subtraction weights are defined (the background-subtracted full data sample has an estimated contamination of 1% from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events, see Section 6.8.3). The fit of the nominal model to the data in

the restricted B mass range, without background subtraction weights, yields a blinded polFrac value of 0.913 ± 0.010 . The magnitudes and phases and the fit fractions of the amplitudes are listed in Table 6.10 and the projections of the fitted model on the five visualisation variables are shown in Figure 6.8.

Following the reasoning from Section 6.8.3, specifically, the toy studies on the effect of backgrounds on the fitted polarisation parameter, the fitted polFrac value should be corrected for the presence of backgrounds in the upper half peak samples. These studies were performed specifically for the π^0 background, but any component of the data other than the signal would lack the specific polarisation structure of the signal model that is fitted and therefore “wash out” the polarisation measurement, *i.e.* bias it in the direction of no polarisation (polFrac value 0.5). Let us note here that since the measured blind values of polFrac are above 0.9, and the blinding offset is between -0.2 and 0.2 , we know that the measured values are greater than 0.5 and therefore the corrections must be positive in order to increase the degree of polarisation. With a correction of 3.7% on the upper half peak result, we get a blind polFrac value of 0.950. With statistical uncertainties only, the blinded polFrac values after correction for background contaminations are 0.945 ± 0.007 for the result from data with background-subtraction weights (see Section 6.8.3) and 0.950 ± 0.010 for the result from the upper half peak data without background subtraction weights. The two samples are statistically dependent: all events in the upper half peak samples are included, with background-subtraction weights, in the nominal samples. Assuming full correlation, the two uncertainties are subtracted in quadrature to determine the uncertainty of the difference between the two polFrac values. The actual uncertainty of the difference would be slightly higher than this simple estimate due to the presence of weights in one of the samples. With the simple estimate of full correlation, the difference between the two values is 0.005 ± 0.007 , which is compatible with zero. This means that there is no indication of bias on the photon polarisation measurement from the background subtraction method. Therefore, no systematic uncertainty related to the background subtraction method will be assigned to the measurement of λ_γ . Since the estimate of the background contamination in the upper half peak sample comes from the same B mass fit as the shapes that are used to define the background-subtraction weights, systematic effects in the fit (*e.g.* the description of the background and signal shapes) could bias the measurement in a way that escapes detection through this cross-check. That is why we will determine a systematic uncertainty related to the B mass fit.

Table 6.10 – Results for the magnitudes and phases and the fit fractions of the amplitudes from the model in Table 6.9, fitted to unweighted data samples with tightened cuts on the B candidate mass. The fit fractions in the rightmost column refer to the decay of the particle that is at the head of the decay in the “Amplitude” column on the same line. For particles with more than one decay channel, the sum of the fit fractions is given.

polFrac (blind)	0.913 ± 0.010		
Amplitude k	Magnitude a_k	Phase ϕ_k (rad)	Fraction (%)
$B^+ \rightarrow K_1(1270)^+ \gamma$	1.000 (fixed)	0.000 (fixed)	39.3 ± 1.3
$K_1(1270)^+ \rightarrow K^+ \rho(770)^0$	1.326 ± 0.067	5.415 ± 0.057	39.2 ± 2.2
$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	21.0 ± 1.3
$K_1(1270)^+ \rightarrow K_0^*(1430)^0 \pi^+$	0.374 ± 0.030	3.377 ± 0.077	6.5 ± 0.9
$K_1(1270)^+ \rightarrow K^*(892)^0 \pi^+$ [D]	2.607 ± 0.160	2.464 ± 0.068	5.0 ± 0.5
Sum($K_1(1270)^+$)			71.7 ± 1.4
$B^+ \rightarrow K_2^*(1430)^+ \gamma$	1.858 ± 0.083	1.397 ± 0.060	8.7 ± 0.4
$K_2^*(1430)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	67.6 ± 2.3
$K_2^*(1430)^+ \rightarrow K^+ \rho(770)^0$	0.576 ± 0.037	0.050 ± 0.069	20.6 ± 1.9
Sum($K_2^*(1430)^+$)			88.1 ± 0.4
$B^+ \rightarrow K(\text{NR}, 1^-)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	0.762 ± 0.058	1.310 ± 0.087	6.6 ± 0.9
$B^+ \rightarrow K^*(1680)^+ \gamma$	0.518 ± 0.030	4.454 ± 0.097	6.6 ± 0.6
$K^*(1680)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	44.7 ± 4.9
$K^*(1680)^+ \rightarrow K^+ \rho(770)^0$	0.728 ± 0.089	4.091 ± 0.147	34.5 ± 4.7
Sum($K^*(1680)^+$)			79.3 ± 0.9
$B^+ \rightarrow K_1(1400)^+ \gamma$	0.350 ± 0.032	5.353 ± 0.091	4.3 ± 0.5
$K_1(1400)^+ \rightarrow K^+ \rho(770)^0$	1.056 ± 0.143	0.937 ± 0.134	55.8 ± 5.8
$K_1(1400)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	31.0 ± 5.3
Sum($K_1(1400)^+$)			86.8 ± 0.9
$B^+ \rightarrow K(\text{NR}, 1^+)^+ \gamma$	0.481 ± 0.041	1.180 ± 0.109	4.3 ± 0.6
$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0$	1.000 (fixed)	0.000 (fixed)	93.7 ± 2.0
$K(\text{NR}, 1^+)^+ \rightarrow K^+ \rho(770)^0$ [D]	0.733 ± 0.106	0.854 ± 0.148	7.5 ± 2.0
Sum($K(\text{NR}, 1^+)^+$)			101.3 ± 0.7
$B^+ \rightarrow K_2(1770)^+ (\rightarrow K^+ \rho(770)^0) \gamma$	0.871 ± 0.064	2.219 ± 0.095	1.4 ± 0.2
$B^+ \rightarrow K^*(1410)^+ \gamma$	0.454 ± 0.039	3.927 ± 0.085	1.4 ± 0.2
$K^*(1410)^+ \rightarrow K^*(892)^0 \pi^+$	1.000 (fixed)	0.000 (fixed)	108.8 ± 2.5
$K^*(1410)^+ \rightarrow K^+ \rho(770)^0$	0.399 ± 0.112	0.735 ± 0.249	17.0 ± 9.1
Sum($K^*(1410)^+$)			125.8 ± 6.8
Sum(B^+)			72.6 ± 1.5

Chapter 6. Amplitude analysis




Figure 6.8 – Projections on the five visualisation variables (see Section 6.2) of the data (black points) and MC integration samples reweighted to the best fit amplitude model (pink) and the sub-models for the different kaonic resonances (colours indicated in the plot legends) in the restricted B mass range from 5300 to 5600 MeV/c² (“upper half peak” samples). The pulls (see Section 6.2) are shown below each histogram.

6.8 Systematic uncertainties

In this section, the main systematic uncertainties that affect the measurement of λ_γ are discussed and quantified in terms of uncertainties on polFrac, which must be doubled to obtain the uncertainties on λ_γ . The largest systematic uncertainty arises from the contamination of the background-subtracted data sample from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events. It is discussed in Section 6.8.3. Further systematic uncertainties related to the momentum resolution (Section 6.8.1), the fixed parameters in the B mass fit described in Chapter 5 (Section 6.8.2), the size of the MC integration samples (Section 6.8.4), the choice of decay amplitudes in the nominal model (Section 6.8.5), and the parametrisation of the $\rho^0 - \omega^0$ resonance mixing (Section 6.8.6) are also discussed. Systematic effects related to the fixed parameters in the amplitude fit are expected to contribute to the systematic uncertainty on the measurement of λ_γ at a secondary level. The strategy for their quantification is discussed (Sections 6.8.7), but the results of the studies are not described in this work. Instead, an estimate based on a simpler study is given. As explained in Section 6.7, no systematic uncertainty related to the background subtraction method itself is assigned to the measurement of λ_γ . The uncertainties on both polFrac and λ_γ are collected in Table 6.11.

Table 6.11 – Systematic and statistical uncertainties (and their quadratic sum) on the photon polarisation parameter both in terms of $\text{polFrac} = \frac{1+\lambda_\gamma}{2}$ and λ_γ . The corrections on the parameters for the effect of the π^0 background in the otherwise background-subtracted data sample are also given, but not included in the quadratic sums of the uncertainties.

	polFrac	λ_γ
Momentum resolution	0.0025	0.0050
Fixed parameters in B mass fit	0.0023	0.0047
π^0 background correction	+0.0116	+0.0232
π^0 background uncertainty	0.0030	0.0060
Integration sample size	0.0017	0.0034
Choice of decay amplitudes	0.0024	0.0048
$\rho^0 - \omega^0$ line shape	0.0024	0.0048
Fixed parameters in amplitude fit (estimate)	0.0016	0.0031
Total systematic uncertainty	0.0061	0.0123
Statistical uncertainty	0.0070	0.0140
Total uncertainty	0.0093	0.0186

6.8.1 Momentum resolution

To determine the systematic effect of the use of reconstructed, rather than true, four-momenta in the fit on data, the “norm” MC samples are used. The samples for each data-taking year are split into sub-samples with a size approximately corresponding to the statistical power of the data sample for that year. This way, nine exclusive sub-samples of the “norm” MC samples are obtained. These nine samples are re-weighted to match the nominal model from the fit on data with an arbitrarily chosen value of 0.885 for polFrac, since the polarisation is blinded in the fit on data. Each MC event is assigned a weight that is the ratio of the PDF of the nominal model with $\text{polFrac} = 0.885$ to the generation PDF value of the “norm” model. No PID weights (see Section 5.1) are applied to any samples in these studies, since the MC samples used here only contain signal events where the final-state particles are correctly identified in the reconstruction. The nine samples are fitted with the nominal data model. The masses and widths of the resonances are free parameters with Gaussian constraints like in the model building. The magnitudes and phases of the amplitudes, as well as polFrac, are also free. The reconstructed, B -mass-constrained (see Section 6.3) momenta are used for the re-weighted MC sub-samples in one set of fits, while the true momenta are used in the other. In all fits, the full “norm” MC samples are used as integration samples.

The mean value of the fitted polFrac parameter over the nine samples is 0.8854 ± 0.0017 for the fits using the true momenta and 0.8825 ± 0.0018 for the fits using the reconstructed, B -mass-constrained momenta. The uncertainties on these values are the square roots of the variances of the nine fitted polFrac central values for each respective set of fits. The mean of the polFrac values fitted with true momenta is compatible with the value of 0.885 from the model used in the re-weight, while the mean from the fits that use B -mass-constrained momenta is 0.0029 lower than the fit result from true momenta and 0.0025 lower than 0.885. Hence, a systematic uncertainty of 0.0025 related to the momentum resolution is assigned to the measured value of polFrac.

Analogous fits are also run using the reconstructed momenta without the B mass constraint. Their mean fitted polFrac value is 0.8816 ± 0.0017 . While this result is not used in the calculation of the systematic uncertainty, it is a useful cross-check of the B mass constraint method. The fact that the fitted value of polFrac for the B -mass-constrained momenta is closer to the true value than for the unconstrained reconstructed momenta confirms that the B -mass-constrained momenta should be used in the fit on data.

6.8.2 Fixed parameters in the B mass fit

The parameterisation of the signal and backgrounds as well as other choices made in the B mass fit (see Section 5.2) can influence the measurement of λ_γ , specifically through their effects on the background subtraction. These effects are estimated by varying the B mass fit

and noting the changes, if any, in the measured value of λ_γ . Two variations of the B mass fit are studied:

- An exponential model is used for the combinatorial background instead of a linear model. The fitted polFrac value changes by 8.8×10^{-5} .
- The relative yield of the $B^+ \rightarrow K\pi\pi\eta$ background to the signal yield is varied by changing the contamination of that background by one unit of its uncertainty. Increasing (decreasing) the η contamination in the B mass fit changes the fitted polFrac value by $+0.0022$ (-0.0021).

The changes in η background yields in the COWs range between $5000 \text{ MeV}/c^2$ and $6500 \text{ MeV}/c^2$ from the two studies are about evenly compensated by the signal yield and the yield of the partially reconstructed background with one missing pion. The results of this study hint, at first sight, at an underestimated background yield, but the effect is below the precision of the “upper half peak” cross-check in Section 6.7.

The fixed shape parameters in the B mass fit other than the η contamination and the π^0 background shape parameters are

- the four tail-shape parameters of the signal CB (see Eq. 5.1);
- two shape parameters each for the Argus functions (Eq. 5.2) of the η background and the partially reconstructed background with one missing pion;
- one scale factor for the width of the resolution function of the partially reconstructed background with one missing pion.

Note that we ignore the fixed contamination and shape parameters of the π^0 background for this systematic study, since a separate systematic effect related to this background is determined in Section 6.8.3. One could also consider the end-points of the three partially reconstructed backgrounds, which are pegged to the signal mean minus zero (for η), one (for one missing pion), and two (for the generic partially reconstructed background) times the π^0 mass, as fixed parameters, but they are not determined from fits to MC samples but from phenomenological considerations. The impact of the nine fixed shape parameters above on the COWs, and therefore on the polarisation measurement, is expected to be much smaller than that of the η contamination. This is because the latter changes the relative yields of the signal and background species, while the other parameters change the shapes, but are expected to have only small effects on the yields. The study of the combinatorial background above also changes the shape of one of the background species, so its result can be used to estimate the effect of the remaining fixed parameters. We multiply the result of the combinatorial background study by ten before adding it, in quadrature, to the η contamination study. The result of 0.0023 is assigned as a systematic uncertainty to the polFrac parameter.

6.8.3 Contamination from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events

The IsPhoton classifier gives values close to 1 for photon-like calorimeter objects, and values closer to 0 for π^0 -like calorimeter objects. By applying the offline selection excluding the requirement on IsPhoton, one can study the distribution of the IsPhoton variable in data samples and compare it to the true γ distribution and true π^0 distribution.

Assuming that the photon candidates in the $B \rightarrow K\pi\gamma$ data sample come from either real photons or real π^0 , the distribution of the IsPhoton variable is understood as being the sum of these two components: for each bin i , the number of events in data $N_{\text{data}}(i)$ is the sum of the contribution from photons and the contribution from π^0 in that bin:

$$N_{\text{data}}(i) = N_{\pi^0} n_{\pi^0}(i) + N_{\gamma} n_{\gamma}(i), \quad (6.11)$$

where N_{π^0} (N_{γ}) is the total number of true π^0 (γ) and $n_{\pi^0}(i)$ ($n_{\gamma}(i)$) the fraction of that number in bin i . When the distributions are known, a binned χ^2 fit can be used to estimate N_{γ} and N_{π^0} .

The IsPhoton classifier is not well modeled in simulation, so the cut efficiencies have to be calibrated using distributions of IsPhoton from pure (background-subtracted) data samples of photons and π^0 as described in Section 5.1. These same calibration samples are used for this study to obtain the distributions n_{γ} and n_{π^0} of the IsPhoton variable for true γ and true π^0 . The π^0 sample comes from 2017 data that has been selected to contain almost exclusively $D^{*+} \rightarrow D^0 (\rightarrow K^+ \pi^- \pi^0) \pi^+$ decays [74]. The subtraction of the remaining background is based on the difference between the D^{*+} and the D^0 candidate invariant masses. The photon sample comes from 2017 data that has been selected for $B^0 \rightarrow K^{*0} \gamma$ decays. The background subtraction is based on the B^0 candidate invariant mass [74]. Since the ECAL is split into three regions with different cell sizes and occupancies (see Section 3), the IsPhoton variable is actually composed of three separately trained classifiers, one for each ECAL region [47]. The distributions in each region should therefore be studied separately.

The histograms of the IsPhoton score $N_{\text{data}}(i)$ are obtained for each ECAL region from data recorded in 2017, the same year as the calibration samples, with the full selection described in Section 5.1 applied except for the requirement on IsPhoton, and with background-subtraction weights as calculated from the distributions of that sample in the B mass, as explained in Section 5.3. In the background subtraction weighting, the signal and $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ shapes are joined, therefore, the weighted sample should include the background from π^0 .

The discrimination power as well as the distributions of the IsPhoton classifiers can depend on the kinematic attributes of the photon candidate. The true distributions of γ and π^0 should therefore be obtained from samples that kinematically match the data they are fitted to. To this end, the calibration samples are weighted in each ECAL region so that their p_T and η distributions match those of the data sample in that ECAL region. These weights are applied in addition to the background-subtraction weights by multiplication. To obtain these weights,




Figure 6.10 – Binning used in the kinematic reweighting of the π^0 calibration sample in the inner region of the ECAL. All histograms have the π^0 candidate p_T [MeV/c] on the horizontal axis and η on the vertical axis. The left column shows the finely binned distributions of the data (top) and the π^0 calibration sample (bottom) in the inner region of the ECAL. The middle column shows the same samples, in a coarser binning that was used to determine the weights. The bottom histogram on the right shows the ratio of the normalised histogram on the top of the middle column to the normalised histogram on the bottom of the middle column; it defines the weights that are assigned to the events in the π^0 calibration sample in the inner region of the ECAL.

the 2D $p_T - \eta$ distributions are binned in each sample and each event in the calibration samples is assigned a weight that is proportional to the ratio of the number of data events in that bin to the number of calibration events in that bin. As an example, the bins in the inner region of the ECAL with data from this analysis as well as from the π^0 calibration sample are shown in Figure 6.10.

N_{π^0} and N_γ are determined in each region of the ECAL by minimising the discrepancy between the data histogram and the scaled γ and π^0 distributions. The discrepancy is quantified by the χ^2 test statistic: for each bin i , the residual r_i is squared and divided by its variance $V(r_i)$, and the sum of these scaled squared residuals gives the χ^2 :

$$\chi^2 = \sum_i \frac{r_i^2}{V(r_i)}. \quad (6.12)$$

The residual r_i is defined as

$$r_i = N_{\text{data}}(i) - (N_{\pi^0} n_{\pi^0}(i) + N_\gamma n_\gamma(i)). \quad (6.13)$$

The variance $V(r_i)$ is calculated from the data and the calibration samples. Since the samples are weighted, the number of events in each bin i is in fact the sum of the weights of those events in the sample that fall into bin i . For example,

$$N_{\text{data}}(i) = \sum_{j \in i} w_j \quad (6.14)$$

and

$$V(N_{\text{data}}(i)) = \sum_{j \in i} w_j^2. \quad (6.15)$$

For the uncertainties on the binned π^0 and γ distributions, the normalisation must be taken into account as well:

$$V(n_{\pi^0}(i)) = \frac{V(N_{\pi^0}(i))}{(\sum_j N_{\pi^0}(j))^2}, \quad (6.16)$$

where $N_{\pi^0}(i) = \sum_{j \in i} w_j$ and $V(N_{\pi^0}(i)) = \sum_{j \in i} w_j^2$ (note that $n_{\pi^0}(i) = \frac{N_{\pi^0}(i)}{\sum_j N_{\pi^0}(j)}$). The variance $V(n_{\gamma}(i))$ is calculated analogously to $V(n_{\pi^0}(i))$. The variance for the residual of bin i is then

$$V(r_i) = V(N_{\text{data}}(i)) + (N_{\pi^0})^2 V(n_{\pi^0}(i)) + (N_{\gamma})^2 V(n_{\gamma}(i)). \quad (6.17)$$

The fit results are shown in Figure 6.11.

Without a requirement on IsPhoton, the overall π^0 contamination is found to be $\frac{N_{\pi^0}}{N_{\gamma}} = (6.3 \pm 0.5)\%$, combining all ECAL regions. From the reweighted π^0 and γ distributions of IsPhoton, the efficiencies of cuts on IsPhoton for π^0 and γ can be evaluated. They are plotted in Figure 6.12, while selected values are tabulated in Table 6.12. For the nominal requirement of the analysis, $\text{IsPhoton} > 0.9$, the contamination is only $(0.86 \pm 0.07)\%$. The $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ contamination values that are calculated using the selection efficiencies and branching fractions, see Table 5.8 in Section 5.2, are $(1.31 \pm 0.21)\%$ and $(1.46 \pm 0.23)\%$ for Run 1 and Run 2, respectively. The two methods of calculating the contamination are very different, though somewhat correlated by the use of the same calibration samples (they are used to determine the efficiency of the IsPhoton cut on MC, thus contributing to the estimates of the selection efficiencies), and consequently are a useful cross-check. Both result in contaminations of the order of 1%; the fact that the values are not compatible indicates that the uncertainty of at least one of the two methods is underestimated.

The effect of a contamination from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ in the signal-only sample on the photon polarisation parameter is estimated from fast simulation studies. A simulated signal sample with 30 000 events is merged with a smaller background sample. To save time, the simulated decays in these samples are not passed through the full LHCb reconstruction and selection process. They are generated in the B rest frame and no selection is applied. The four-momentum values of the π^0 in the background sample are assigned to the γ candidate in the fit, and the resulting sample is fitted with a signal model. Since the π^0 is a pseudoscalar, a fit to a background sample should not indicate any polarisation, and indeed, fits to background samples

6.8. Systematic uncertainties




Figure 6.11 – Fits to the IsPhoton distribution in background-subtracted 2017 data after the offline selection excluding the photon PID cuts in linear (top) and logarithmic (bottom) scale. The fits are done separately in the three regions of the ECAL: the inner region (Region 1, left), the middle region (Region 2, centre), and the outer region (Region 3, right).




Figure 6.12 – IsPhoton cut efficiencies for the γ (blue) and π^0 (orange) calibration samples, reweighted to match the data's $p_T - \eta$ distributions, in the three regions of the ECAL: inner (left), middle (centre), and outer (right). The nominal requirement is $\text{IsPhoton} > 0.9$.

Table 6.12 – Efficiency of photons and π^0 in background-subtracted 2017 data sample at different IsPhoton cut values. The contamination, defined as the ratio of π^0 to γ events, is estimated from a fit to the IsPhoton distribution. Note that requiring IsPhoton to be greater than zero is not the same as not placing a requirement on IsPhoton as the distributions extend slightly into negative values.

cut	γ efficiency (%)				π^0 efficiency (%)				π^0 contamination (%)			
	ECAL region				ECAL region				ECAL region			
	1	2	3	all	1	2	3	all	1	2	3	all
0.00	100.0	100.0	100.1	100.0	96.4	89.8	95.8	94.9	8.7	3.2	6.2	6.0
0.38	95.1	97.7	99.1	97.4	55.0	54.5	65.7	59.0	5.2	2.0	4.3	3.8
0.70	83.7	90.8	92.8	89.4	27.4	27.2	34.4	30.0	2.9	1.1	2.4	2.1
0.86	65.2	78.9	79.3	75.0	12.4	13.5	17.1	14.4	1.7	0.6	1.4	1.2
0.90	53.7	70.2	67.6	64.3	7.7	9.3	9.9	8.8	1.3	0.5	1.0	0.9
0.95	23.5	41.4	34.3	33.4	2.1	3.2	2.3	2.4	0.8	0.3	0.4	0.5

yield values of the polarisation fraction compatible with 0.5 (0.491 ± 0.011 and 0.511 ± 0.027 from fits to ten simulated $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ samples with model C (see below) with “val” and “norm” fit models respectively) which corresponds to no polarisation. The addition of an unpolarised-like component to a polarised signal sample should dilute the polarisation of the sample: the bias on the fitted polarisation is expected to be roughly proportional to the relative amount of background events in the sample. As an example, when a fully polarised signal sample and an unpolarised signal sample (unpolarised for the signal meaning that 50% of the photons are right-handed and 50% are left-handed) of the same size are merged, the resulting sample is 75% polarised, and a fit to the merged sample should return a polarisation fraction of 0.75 (or a polarisation fraction of 0.25, depending whether the polarised sample is right-handed or left-handed). Adding an unpolarised (polarisation fraction 0.5) background component to a signal sample with a polarisation fraction of 0 should therefore yield a fitted polarisation fraction of $R/2$, where R is the proportion of background events in the combined sample.

The bias on the polarisation fraction may also depend on the make-up of the π^0 background (*i.e.* the π^0 background model), or on the signal model that is used in the fit.

To test the dependence of the bias on the background model and the fit model, a signal sample with 30 000 events is generated using the “val” signal model. Background samples of $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ decays are also generated with four different models:

- Model A: all the decay chains of the “val” model, but swapping the γ for a π^0 ;
- Model B: same as model A, but with two extra amplitudes involving the pseudo-scalar resonance $K(1460)^+$, which is forbidden in decays with a photon;
- Model C: same as model B, but with different phases for the pseudo-scalar amplitudes;

- Model D: only the two pseudo-scalar amplitudes, plus a third non-resonant scalar amplitude.

For each background sample and model, the signal sample is merged with 5000 (1000, 300) events from the background sample, yielding a 35000 (31000, 30300) event sample which has a relative proportion of background events of 14.2% (3.2%, 1%). The samples are fitted with the “val” model and with the “norm” model in turn, to cover both the case where the true signal model is assigned, and the case where the assigned model has more degrees of freedom than the signal model which could be picked up by the background contribution. The results are summarised in Table 6.13.

Table 6.13 – Results of studies on the effect of a π^0 background contamination on the fitted photon polarisation parameter for different π^0 background models (see text) and different proportions R of π^0 background events in the full sample. The signal is generated with a true polFrac of 0.

π^0 bkg model	R	Fit model	Fitted polFrac
A	14.2%	val	0.125 ± 0.006
B	14.2%	val	0.120 ± 0.006
C	14.2%	val	0.119 ± 0.006
D	14.2%	val	0.113 ± 0.006
A	14.2%	norm	0.124 ± 0.006
B	14.2%	norm	0.120 ± 0.006
C	14.2%	norm	0.117 ± 0.006
D	14.2%	norm	0.116 ± 0.006
A	3.2%	val	0.035 ± 0.006
B	3.2%	val	0.035 ± 0.006
C	3.2%	val	0.034 ± 0.006
D	3.2%	val	0.028 ± 0.006
A	3.2%	norm	0.036 ± 0.006
B	3.2%	norm	0.036 ± 0.006
C	3.2%	norm	0.032 ± 0.006
D	3.2%	norm	0.031 ± 0.006
A	1%	val	0.012 ± 0.005
B	1%	val	0.011 ± 0.005
C	1%	val	0.011 ± 0.005
D	1%	val	0.011 ± 0.005
A	1%	norm	0.010 ± 0.005
B	1%	norm	0.010 ± 0.005
C	1%	norm	0.009 ± 0.005
D	1%	norm	0.009 ± 0.005

The bias on the polarisation fraction introduced by the background events is of the same order as the proportion of background. This is the case both for the fit model that matches

the generated signal model (“val”) and for the fit model that does not (“norm”). This result is also independent of the precise content of the background model. The bias is systematically larger than the expectation of $R/2$ for an unpolarised background component. To further study the bias, the study is extended for one of the π^0 background models (model C). Relative proportions R of 1%, 3.2%, 14.2%, 33.3%, 42.9%, and 50% are tested, corresponding to 300, 1 000, 5 000, 15 000, 22 500, and 30 000 background events added to the signal sample of 30 000 events. The previous signal sample, generated with the “val” model with a polarisation fraction of 0, is tested, as well as a signal sample generated with the “val” model but a polarisation fraction of 1. In addition to testing the effect of a background contamination, we also test the effect of a contamination with unpolarised signal events. To this end, a signal sample is generated with the “val” model with a polarisation fraction of 0.5 and used to contaminate the usual polarised “val” signal sample as well as a signal sample generated with the “norm” model with a polarisation fraction of 0. The fitted values of the polarisation fraction for different combinations of polarised signal and background or unpolarised signal samples are shown in Figure 6.13. The results from samples with unpolarised signal contamination broadly match the expectation of $R/2$, but the biases from contamination with π^0 background are systematically larger. In the region of low contamination values, however, the discrepancy is small, and additionally, the biases as well as their differences are of the same order as the uncertainty on the parameter from the fit.

The correction of the polarisation parameter for the presence of π^0 background in the background-subtracted data sample is applied using the result from background contamination rather than signal contamination. The contamination being around 1%, a correction equal to the contamination is applied, rather than a correction of half the contamination.

The first method to estimate the π^0 contamination uses selection efficiencies and branching fractions and leads to estimates of the contamination of $(1.31 \pm 0.21)\%$ in Run 1 and $(1.46 \pm 0.23)\%$ in Run 2. The second method, using the distribution of the IsPhoton variable, gives a contamination of $(0.86 \pm 0.07)\%$ for 2017 data. The two methods are completely independent of each other. Both indicate that the measured polarisation parameter should have a correction on the order of 1% applied. The fact that they are only compatible within 2.5σ must be reflected in the systematic uncertainty associated with the π^0 background. We choose the following prescription for the correction and the systematic uncertainty: since the estimate from the IsPhoton distribution concerns data from 2017, the resulting value is compared to the value from the branching fractions for Run 2. The different signal and background selection efficiencies between Run 1 and Run 2 do not change the contamination to such an extent that the IsPhoton study needs to be repeated for Run 1 data, as is evidenced by the compatible contamination estimates from the first method. The correction and systematic uncertainty is therefore applied to the global (using data from Run 1 and Run 2) result for the polarisation parameter. The correction applied is equal to the arithmetic average of the two contamination estimates: 1.16%. With a blinded polFrac of 0.933 in the nominal fit, we get a corrected value of 0.945. The two estimates are 0.6% apart, so we assign a systematic uncertainty of half this difference, 0.3%, to the measured polFrac value in addition to the correction.




Figure 6.13 – Fitted values of the polarisation fraction (polFrac) for different combinations of polarised signal samples and background (circles) or unpolarised signal (crosses) samples. The expected bias of $R/2$ is shown in black. For the signal sample with a polarisation fraction of 1, the plot shows $1 - \text{polFrac}$ rather than polFrac. All the contamination values are shown in the top plot, while the bottom plot shows a zoomed view of the results for low contamination values, with a meaningless offset in the abscissa to stagger the points for visibility and with the statistical uncertainty of the fit results as error bars.

6.8.4 Integration sample size

The integration samples must contain enough events to precisely compute the PDF integral everywhere in the phase space of the signal model so as not to bias the measurement. To ensure this, very large MC samples with signal-like properties were produced (see Section 6.3). Any remaining bias is therefore expected to be small. To determine the associated systematic uncertainty, the integration sample is split into three sub-samples of approximately equal size, and the fit to the nominal model is repeated, each time with the full set of background-subtracted data events. The largest difference between the three results for the polFrac parameter is divided by $\sqrt{3}$ and the result of 0.0017 is assigned as a systematic uncertainty.

6.8.5 Choice of decay amplitudes

The large number of amplitudes in Tables 6.4 and 6.5, and the even larger number of possible combinations of them, makes it impractical to test all possible sets of decay amplitudes to find the model that best fits the data, which is why the nominal model is found through an iterative algorithm. In order to quantify the effect that a different choice of decay amplitudes would have on the measurement of λ_γ , alternative decay models must be tested. To choose these alternative models arbitrarily, however, would discard the optimised description of the data that is the result of the model building, and therefore would over-estimate the systematic uncertainty associated with model content choice. Instead, two alternative models are tested: for one, the amplitude with the lowest significance, which is also the amplitude with the lowest fit fraction, that of the decay $B^+ \rightarrow K^*(1410)^+ (\rightarrow \rho^0 (\rightarrow \pi^+, \pi^-), K^+) \gamma$ (see Section 6.6), is removed from the nominal model. The other alternative model is obtained by adding the amplitude for the decay $B^+ \rightarrow K_1(1270)^+ (\rightarrow K^+ \rho^0 (\rightarrow \pi^+ \pi^-))$ [D-wave] γ from step 6 of the iterative part of the model building (see Table 6.7) to the nominal model. Both alternative models are fitted to the data. The former alternative model, with one fewer amplitude than the nominal, gives a polFrac value that differs more from the nominal result than the latter, and this larger of the two differences of 0.0024 is assigned as a systematic uncertainty.

6.8.6 The parametrisation of the $\rho^0 - \omega^0$ resonance mixing

The parametrisation of the main resonances in the $\pi^+ \pi^-$ system used in the nominal model is a Gounaris-Sakurai (GS) line-shape for the $\rho(770)^0$, which is adapted to include the $\omega(782)^0$ through a multiplication by a factor of $(1 + \delta \frac{s}{m_\omega^2} \text{BW}_\omega)$, where BW_ω is a Breit-Wigner (BW) function for the $\omega(782)^0$ resonance, and the parameter δ is set to 0.00471 (see Section A.2.3 in Appendix A.2). This parametrisation leads to a good description of the data. The fact that it strictly constrains the contribution from the $\omega(782)^0$ to that of the $\rho(770)^0$ is advantageous in the model building and generally in the amplitude fit, as it reduces the fit complexity. A systematic effect from this strong assumption is determined using a different parametrisation

where the ρ is still described by a GS function, and the ω by a BW function, but the two are decoupled. This adds nine decay amplitudes to the model, amounting to 18 additional free parameters in the minimisation. To stabilise the fit, the masses and widths of the two light-flavour resonances are fixed to their PDG values (the K_{res} masses are fixed to the nominal model results); it is also necessary to constrain the relative magnitude of the $\omega(782)^0$ contribution in the decays of vector K_{res} to 10% (the approximate result of some fits whose minima could not be verified due to the strong correlation between the $\rho(770)^0$ and $\omega(782)^0$ amplitudes) of that of the $\rho(770)^0$. The fit to data with this alternative parametrisation yields a blinded polFrac value of 0.931 ± 0.007 for a χ^2/ndf of 1.21. This differs from the nominal blind polFrac by 0.0024. The difference is assigned as a systematic uncertainty.

As a further variation, the nominal description of the $\rho(770)^0$ line-shape is used, but the mixing parameter δ , which is fixed in the nominal model and during model building (see Section A.2.3 in Appendix A.2), is freed and no longer constrained to be real, *i.e.* both the magnitude and the phase of δ are fitted. The masses and widths of the $\rho(770)^0$ and the $\omega(782)^0$ are also free in the fit, with Gaussian constraints as in the model building. The K_{res} masses are fixed. The fit yields 0.0027 ± 0.0002 for the magnitude and -0.098 ± 0.086 for the phase of δ . The fitted values of the $\rho(770)^0$ and the $\omega(782)^0$ masses and widths are compatible with the PDG values, except for the $\omega(782)^0$ width, which is $0.0092 \pm 0.0004 \text{ GeV}/c^2$ (PDG: $0.00868 \pm 0.00013 \text{ GeV}/c^2$). While neither the magnitude nor the phase of δ from the fit are compatible with the value of 0.00471 used in the analysis, the blind polFrac value from that fit is exactly the same as the nominal value. Therefore, no additional systematic uncertainty is assigned.

6.8.7 Fixed parameters in the amplitude fit

The masses of the K_{res} and the masses and widths of the $\rho(770)^0$ and the $\omega(782)^0$ are optimised in fits on data as part of the model building. Their fitted values can be assumed to depend on the parametrisation of the decay amplitudes. To determine a systematic uncertainty related to these parameters and to the decay radii of the resonances in the nominal model, the parameter values are varied and the fit is repeated. An upper bound of the resulting systematic uncertainty can be estimated from a fit with all the masses and widths set to their PDG values rather than to the values from the nominal model. This fit is mentioned in Section 6.6 as a cross-check. It yields a blind polFrac of 0.935 ± 0.007 , which is 0.0016 away from the nominal fit result, so the systematic uncertainty related to the fixed parameters in the amplitude fit is expected to be of the order of 0.0016 or lower.

Chapter 7

Conclusion

A map of the magnetic field of the LHCb detector for Run 3 has been built from simulated and measured field data. This map is created from scratch, where previous updates of the field maps have been based on corrections to older maps, and several improvements have been made to the simulation model, the translation of the simulated field to a map in the LHCb format, the measurement data, and its incorporation into the final map. The new map, with its associated alignment, has been shown to outperform its predecessor in accuracy and precision of reconstructed mass peaks in 2024 data. Given the complete overhaul of the process, the small but significant improvements in these metrics are a welcome success and show that the methods outlined in this work not only provide a magnetic field map that is fit for use in LHCb analyses, but also pave the way towards further refinements of the map. A closer correspondence of the field map to the real field could be achieved with, for example, field measurements beyond the region of high field values that covers the measurements used in this work, and/or a more detailed knowledge of the ferromagnetic structures in the LHCb cavern, specifically the reinforced concrete in the cavern floor and walls.

A measurement of the photon polarisation in $b \rightarrow s\gamma$ transitions is performed through an amplitude analysis of $B^+ \rightarrow K^+\pi^-\pi^+\gamma$ decays, using data recorded at the LHCb experiment in 2011, 2012, 2015, 2016, 2017, and 2018. Backgrounds are reduced through tailored selection requirements. The remaining backgrounds are subtracted from the data using weights calculated with COWs based on signal and background distributions in the reconstructed B mass, with the exception of the background from $B^+ \rightarrow K^+\pi^-\pi^+\pi^0$ decays. The mass distributions of the signal and the $B^+ \rightarrow K^+\pi^-\pi^+\pi^0$ background are too similar to allow subtraction of this background. It is estimated from selection efficiencies and branching fraction ratios, as well as from fits to the distributions of a dedicated γ/π^0 discriminant in data, that about 1% of the photon candidates in the background-subtracted data sample are neutral pions that have been misidentified as photons. Instead of modelling the $B^+ \rightarrow K^+\pi^-\pi^+\pi^0$ background, the degree to which it dilutes the measured polarisation is studied using simulation, and the dilution effect is corrected for in the measurement.

The signal decay is described in the isobar formalism. It proceeds through decay chains involving hadronic resonances, where the photon is always produced in the first decay of the chain. There are many hadronic resonances that can contribute, so the number of possible decay chains is high. A selection of decay chains is made in an iterative process called “model building” where decay chains are added to improve the description of the data. The presence of interferences between different decay chains is crucial for the sensitivity of the $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$ decay to the photon polarisation. The selection efficiency is not parametrised but taken into account in the normalisation of the amplitude fit PDF, which uses a large sample of simulated signal events with LHCb reconstruction and full selection applied. The hadronic resonances are mostly described using Breit-Wigner line-shapes, with some resonances, such as the $\rho(770)^0$, having specific line-shapes that have been experimentally established in addition to their motivation from theory. The radiative part of the decay is described with spin factors derived in the helicity formalism. They are added as custom spin factors to the AMPGEN framework, in which the amplitude fit on data is performed. The same framework is used to generate the simulated events used for the PDF normalisation. Both the amplitude fit and the model building are tested and validated on simulated signal samples. The iterative model building on data is deemed successful since the figure of merit that quantifies the agreement between the amplitude model and the data and the polarisation parameter both stabilise.

The resulting nominal amplitude model describes the data well, based on projections of the model and the data on five visualisation variables. The measured polarisation parameter λ_γ , which has a physical range from -1 to $+1$ and is predicted to be close to $+1$ in the SM, is concealed from the researchers at the time of writing, to avoid biasing the measurement. Its statistical uncertainty is 0.0140. A cross-check of the background subtraction method with a data sample where stricter selection requirements than the nominal have been applied, sacrificing some statistical power, indicates no bias on the measurement from the background subtraction method. As a consequence, no systematic uncertainty is assigned to the background subtraction. Other systematic uncertainties on the measured value of λ_γ have been determined for the momentum resolution, the fixed parameters in the B mass fit, the correction of the polarisation dilution from $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ events, the integration sample size, the choice of decay amplitudes, and the parametrisation of the light-flavour resonance line-shape. An additional systematic uncertainty from fixed parameters in the amplitude fit has not been determined in this work. For the purposes of discussing the sensitivity of this measurement, it is estimated at 0.0031 (see Section 6.8.7). The total systematic uncertainty on the measurement of λ_γ is then 0.0123, compared to the statistical uncertainty of 0.0140. With a combined uncertainty of 1.9% on λ_γ , the measurement of the photon polarisation in $b \rightarrow s\gamma$ transitions discussed in this work constrains New Physics contributions to the Wilson coefficients \mathcal{C}_7 and \mathcal{C}'_7 about ten times as strongly as the first direct measurement of the photon polarisation in $b \rightarrow s\gamma$ transitions from $\Lambda_b \rightarrow \Lambda\gamma$ decays in Ref. [14], assuming that the value of λ_γ , once uncovered, is compatible with the SM.

A larger data sample could reduce the uncertainty on λ_γ by lowering the statistical uncertainty, but more importantly, by helping to refine some systematic effects. Specifically, a more

precise estimate of the $B^+ \rightarrow K^+ \pi^- \pi^+ \pi^0$ background contamination, which is one of the main sources of systematic uncertainty, could be obtained with more data, through a more detailed adaptation of the calibration samples' kinematics to the data, for example. The parameters of the nominal model other than λ_γ , including the decay chains that are part of this model, and the fact that some resonances, *e.g.* the $K_2^*(1580)$, have been tested and not found to improve the measurement enough to warrant inclusion in the nominal model, are a secondary result of the amplitude analysis. We found, in the development of the model building, that there is a large number of models with some difference in model content and/or line-shape description that describe the data similarly well. A larger data sample could allow some more qualitative differentiation between these competing descriptions of the hadronic part of the signal decay and increase the sensitivity to the Dalitz-analysis-like part of this amplitude analysis. The polarisation parameter does not differ significantly between the competing models, which lends confidence in the robustness of the method. Nevertheless, a reduction of the model ambiguity could also reduce the systematic uncertainties related to the amplitude model.

The photon polarisation in $b \rightarrow s\gamma$ decays is strongly constrained experimentally, leading to constraints on New Physics contributions to the Wilson coefficients \mathcal{C}_7 and \mathcal{C}'_7 . The measurement discussed in this work, once the remaining systematic uncertainties are quantified and the central value uncovered, will refine these constraints and contribute to lighting the path of extensions to the SM, towards a (more) complete description of elementary and composite particles and their behaviour.

Appendix A

Appendix

A.1 Magnetic field visualisations




Figure A.1 – Differences between the corrected magnetic field measurement values for the B_x component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, before shifting the magnet.

A.1. Magnetic field visualisations




Figure A.2 – Differences between the corrected magnetic field measurement values for the B_z component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, before shifting the magnet.




Figure A.3 – Differences between the corrected magnetic field measurement values for the B_x component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the field map used in LHCb for Run 2.

Appendix A. Appendix




Figure A.4 – Differences between the corrected magnetic field measurement values for the B_y component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the field map used in LHCb for Run 2.




Figure A.5 – Differences between the corrected magnetic field measurement values for the B_z component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the field map used in LHCb for Run 2.

A.1. Magnetic field visualisations




Figure A.6 – Histograms of the field differences (in Tesla) of Figures A.3, A.4, A.5.




Figure A.7 – Differences between the corrected magnetic field measurement values for the B_x component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, after shifting the magnet.

Appendix A. Appendix




Figure A.8 – Differences between the corrected magnetic field measurement values for the B_y component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, after shifting the magnet.




Figure A.9 – Differences between the corrected magnetic field measurement values for the B_z component, taken with “MagDown” polarity in January 2021, and the interpolated values for the same component from the OPERA simulated map with the updated simulation model, after shifting the magnet.

A.1. Magnetic field visualisations




Figure A.10 – Histograms of the field differences (in Tesla) of Figures A.7, A.8, A.9.

A.2 Propagators

Natural units, where $\hbar = c = 1$, are used throughout this section.

A.2.1 Breit-Wigner

The default line shape for a hadronic resonance is the Breit-Wigner propagator [61],

$$\mathcal{T}_{\text{BW}}(s, q, L) = \frac{B_L(q, 0)}{m_0^2 - s - i m_0 \Gamma(s, q, L)}, \quad (\text{A.1})$$

where s is the squared invariant mass of its two daughters, q is the breakup momentum (the absolute value of the three-momentum of one of the daughters in the resonance's rest frame) and L the relative angular momentum of the two daughters, B_L are normalised Blatt-Weisskopf barrier factors, given in Table A.1, m_0 is the nominal mass of the resonance, and constant factors are omitted. The breakup momentum q is a function of s and the masses of the daughters.

The width Γ depends on the breakup momentum q and the relative angular momentum L of the daughters. It is given by

$$\Gamma(s, q, L) = \Gamma_0 \frac{m_0}{\sqrt{s}} \left(\frac{q}{q_0} \right)^{2L+1} B_L(q, q_0)^2, \quad (\text{A.2})$$

where Γ_0 is the width and q_0 the breakup momentum at $s = m_0^2$, for the resonances that decay into two final-state particles.

For the kaonic resonances K_{res} that decay to $K\pi\pi$, the width in the propagators \mathcal{T}_i is computed numerically to account for the dynamics of the different decay chains that lead to the $K\pi\pi$ system as documented in Refs. [65, 76], except for those resonances which have only one decay channel to $K\pi\pi$, for which the running width is given by Eq. A.2.

Table A.1 – Normalised Blatt-Weisskopf centrifugal barrier factors for relative angular momentum L . The radial parameter R used is 1.5 (GeV/c)^{-1} according to a measurement by the Belle collaboration [75].

L	$B_L(q, q_0)$
0	1
1	$\sqrt{\frac{1 + R^2 q_0^2}{1 + R^2 q^2}}$
2	$\sqrt{\frac{9 + 3R^2 q_0^2 + R^4 q_0^4}{9 + 3R^2 q^2 + R^4 q^4}}$

A.2.2 Gounaris-Sakurai mass propagator

The Gounaris-Sakurai line-shape [62] is defined as

$$\mathcal{T}_{\text{GS}}(s, q, L) = \frac{B_L(q, 0)}{m_0^2 - s + \phi(s, q) - i m_0 \Gamma(s, q, L)}, \quad (\text{A.3})$$

where the running width $\Gamma(s, q, L)$ (Eq. A.2) is the same as for the Breit-Wigner propagator. The function $\phi(s, q)$ is

$$\phi(s, q) = \frac{\Gamma_0}{\pi} \left(2 \ln \left(\frac{\sqrt{s} + \sqrt{q^2}}{2m_\pi} \right) \frac{m_0^2 q^3}{\sqrt{s} q_0^3} + \ln \left(\frac{m_0 + \sqrt{q_0^2}}{2m_\pi} \right) \frac{q^2(s - 3m_0^2) + s(m_0^2 - s)}{m_0 q_0^2} + \frac{m_0^2 - s}{q_0} \right). \quad (\text{A.4})$$

A.2.3 A combined propagator for ρ and ω

A single $\pi^+ \pi^-$ line-shape can be used to describe the ρ and ω resonances as well as their mixing. The physical states ρ and ω are mixtures of the isovector ρ_0 and isoscalar ω_0 states. The isovector component dominates the ρ , and the mixing-induced contribution from the ω can be included in the ρ line-shape as [77]

$$\mathcal{T}_{\rho-\omega}(s, q, L) = \mathcal{T}_{\text{GS}_\rho}(s, q, L) \left[1 + \delta \frac{s}{m_{\omega,0}^2} \mathcal{T}_{\text{BW}_\omega}(s, q, L) \right]. \quad (\text{A.5})$$

It uses the Gounaris-Sakurai propagator for the ρ and a Breit-Wigner line-shape for ω . The s -dependence of the expansion is explicit so that the mixing parameter δ is constant. Following the argument in Refs. [78, 79], it is set to $3 \cdot |\delta_{\text{em}}|$, where $|\delta_{\text{em}}| = 0.00157$ [77]. The complex phase of the parameter δ , which accounts for the relative phase between the two resonances, cannot be predicted by an analogous argument. For simplicity, δ is constrained to be real in this analysis. A cross-check in Section 6.8.6 indicates no bias on the polarisation measurement from the choice of value of δ made in this analysis.

A.2.4 The LASS line-shape

Non-resonant S-wave decays to $K^+ \pi^-$ are described by a line-shape derived from scattering experiments by the LASS collaboration [80]. It can be seen as the sum of a non-resonant and a Breit-Wigner propagator,

$$\mathcal{T}_{\text{LASS}}(s, q) = \frac{2a\sqrt{s}}{2 + arq^2 - 2iaq} + \frac{2 + arq^2 + 2iaq}{2 + arq^2 - 2iaq} \mathcal{T}_{\text{BW}}(s, q, 0) \quad (\text{A.6})$$

where the parameters a and r are fixed from the LASS scattering data [81].

A.2.5 The Bugg line-shape

Non-resonant S-wave decays to $\pi^+ \pi^-$ are described by a line-shape based on solutions to the Roy equations [82] and scattering data, proposed in Ref. [83] by D.V. Bugg.

A.3 Elements of the strong decay spin factors

The polarisation tensor for a spin- J boson has rank J . For (pseudo-)scalars, it can be set to 1. Spin-1 particles can have three projections m of their intrinsic angular momentum onto an arbitrary z axis. The polarisation vectors in the particle's rest frame are

$$\epsilon_{(1)}^\mu(m = \pm 1) = \frac{1}{\sqrt{2}} \begin{pmatrix} 0 \\ \mp 1 \\ -i \\ 0 \end{pmatrix}, \quad \epsilon_{(1)}^\mu(m = 0) = \begin{pmatrix} 0 \\ 0 \\ 0 \\ 1 \end{pmatrix}. \quad (\text{A.7})$$

They can be transformed into polarisation vectors $\epsilon_{(1)}^\mu(p, m)$ for spin-1 particles of four-momentum p with Lorentz transformations.

The rank-2 tensors of a spin-2 particle are built from two rank-1 tensors using Clebsch-Gordan coefficients $\langle j_1 m_1, j_2 m_2 | JM \rangle$:

$$\epsilon_{(2)}^{\mu\nu}(p, m) = \sum_{m_1, m_2} \langle 1 m_1, 1 m_2 | 2 m \rangle \epsilon_{(1)}^\mu(p, m_1) \epsilon_{(1)}^\nu(p, m_2). \quad (\text{A.8})$$

Spin projectors are built in their turn from their respective polarisation tensors. For spin-1 particles:

$$P_{(1)}^{\mu\nu}(p) = \sum_m \epsilon_{(1)}^\mu(p, m) \epsilon_{(1)}^\nu(p, m) = -g^{\mu\nu} + \frac{p^\mu p^\nu}{p^2} \quad (\text{A.9})$$

and for spin-2 particles:

$$P_{(2)}^{\mu\nu\rho\sigma} = \sum_m \epsilon_{(2)}^{\mu\nu}(p, m) \epsilon_{(2)}^{\rho\sigma}(p, m). \quad (\text{A.10})$$

The angular momentum tensor for a two-particle state A, B is obtained by projecting the relative angular momentum $q = p_A - p_B$ onto the sub-space of spin L with the appropriate spin projection tensor given above:

$$L_{(L_{AB})}^{\mu_1 \dots \mu_L}(p_A, p_B) = (-1)^{L_{AB}} P_{(L_{AB})}^{\mu_1 \dots \mu_L \nu_1 \dots \nu_L}(p_A, p_B) q_{\nu_1} \dots q_{\nu_L}. \quad (\text{A.11})$$

Bibliography

- [1] W. Haidinger, *Ueber das directe Erkennen des polarisirten Lichts und der Lage der Polarisationsebene*, Annalen der Physik **139** (1844) 29, arXiv:<https://onlinelibrary.wiley.com/doi/pdf/10.1002/andp.18441390903>.
- [2] R. P. O’Shea *et al.*, *Historical Context, Scientific Context, and Translation of Haidinger’s (1844) Discovery of Naked-Eye Visibility of the Polarization of Light*, arXiv:2010.15252.
- [3] Wikimedia Commons, the free media repository, *File:Haidinger's brush.jpg*, 2020. [Online; accessed 16-July-2024].
- [4] G. Horváth *et al.*, *Celestial polarization patterns sufficient for Viking navigation with the naked eye: detectability of Haidinger’s brushes on the sky versus meteorological conditions*, Royal Society Open Science **4** (2017) 160688, arXiv:<https://royalsocietypublishing.org/doi/pdf/10.1098/rsos.160688>.
- [5] Belle collaboration, Y. Ushiroda *et al.*, *New measurement of time-dependent CP-violating asymmetry in $B^0 \rightarrow K_s^0 \pi^0 \gamma$ decay*, Phys. Rev. Lett. **94** (2005) 231601, arXiv:hep-ex/0503008.
- [6] BaBar collaboration, B. Aubert *et al.*, *Measurement of time-dependent CP asymmetry in $B^0 \rightarrow K_s^0 \pi^0 \gamma$ decays*, Phys. Rev. **D78** (2008) 071102, arXiv:0807.3103.
- [7] LHCb collaboration, R. Aaij *et al.*, *First experimental study of photon polarization in radiative B_s^0 decays*, Phys. Rev. Lett. **118** (2017) 021801, arXiv:1609.02032.
- [8] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP-violating and mixing-induced observables in $B_s^0 \rightarrow \phi \gamma$ decays*, Phys. Rev. Lett. **123** (2019) 081802, arXiv:1905.06284.
- [9] LHCb collaboration, R. Aaij *et al.*, *Strong constraints on the $b \rightarrow s \gamma$ photon polarisation from $B^0 \rightarrow K^{*0} e^+ e^-$ decays*, JHEP **12** (2020) 081, arXiv:2010.06011.
- [10] LHCb collaboration, R. Aaij *et al.*, *Observation of photon polarization in the $b \rightarrow s \gamma$ transition*, Phys. Rev. Lett. **112** (2014) 161801, arXiv:1402.6852.
- [11] M. Gronau, Y. Grossman, D. Pirjol, and A. Ryd, *Measuring the photon polarization in $B \rightarrow K \pi \pi \gamma$* , Phys. Rev. Lett. **88** (2002) 051802, arXiv:hep-ph/0107254.

Bibliography

- [12] G. Veneziano, *Towards the measurement of photon polarisation in the decay $B^+ \rightarrow K^+ \pi^- \pi^+ \gamma$* , PhD thesis # 6896, EPFL, 2016.
- [13] V. Bellée *et al.*, *Using an amplitude analysis to measure the photon polarisation in $B \rightarrow K \pi \pi \gamma$ decays*, Eur. Phys. J. **C79** (2019) 622, arXiv:1902.09201.
- [14] LHCb collaboration, R. Aaij *et al.*, *Measurement of the photon polarization in $\Lambda_b^0 \rightarrow \Lambda \gamma$ decays*, Phys. Rev. **D105** (2022) L051104, arXiv:2111.10194.
- [15] C. Elsasser, *Feynman diagram library*, 2014. <http://www.physik.uzh.ch/che/FeynDiagram/index.php>.
- [16] A. J. Buras, *Weak hamiltonian, CP violation and rare decays*, arXiv:hep-ph/9806471.
- [17] K. Chetyrkin, M. Misiak, and M. Münz, *Weak radiative B-meson decay beyond leading logarithms*, Phys. Lett. **B400** (1997) 206, arXiv:hep-ph/9612313.
- [18] M. Gronau and D. Pirjol, *Photon polarization in radiative B decays*, Phys. Rev. **D66** (2002) 054008, arXiv:hep-ph/0205065.
- [19] A. Paul and D. M. Straub, *Constraints on new physics from radiative B decays*, JHEP **04** (2017) 027, arXiv:1608.02556.
- [20] J. Hebinger, *B mesons phenomenology to the search of a signal beyond the Standard Model*, PhD thesis, Université Paris-Saclay, 2017.
- [21] J. M. Blatt and V. F. Weisskopf, *Theoretical nuclear physics*, Springer, New York, 1952.
- [22] A. Einstein, *Ist die Trägheit eines Körpers von seinem Energieinhalt abhängig?*, Annalen der Physik **323** (1905) 639.
- [23] LHCb collaboration, *LHCb magnet: Technical Design Report*, CERN-LHCC-2000-007, 2000.
- [24] LHCb collaboration, A. A. Alves Jr. *et al.*, *The LHCb detector at the LHC*, JINST **3** (2008) S08005.
- [25] J. André *et al.*, *Status of the LHCb dipole magnet*, IEEE Trans. Appl. Supercond. **14** (2004) 509.
- [26] J. André *et al.*, *Status of the LHCb magnet system*, IEEE Trans. Appl. Supercond. **12** (2002) 366.
- [27] M. Losasso *et al.*, *Tests and field map of LHCb dipole magnet*, CERN-LHCb-PROC-2005-028, CERN, Geneva, 2005.
- [28] M. Adinolfi *et al.*, *Performance of the LHCb RICH detector at the LHC*, Eur. Phys. J. **C73** (2013) 2431, arXiv:1211.6759.

- [29] C. Abellan Beteta *et al.*, *Calibration and performance of the LHCb calorimeters in Run 1 and 2 at the LHC*, arXiv:2008.11556, submitted to JINST.
- [30] LHCb collaboration, *LHCb trigger system: Technical Design Report*, CERN-LHCC-2003-031, 2003.
- [31] LHCb collaboration, *Trigger schemes*, <http://lhcb.web.cern.ch/lhcb/speakersbureau/html/TriggerScheme.html>.
- [32] LHCb collaboration, R. Aaij *et al.*, *Design and performance of the LHCb trigger and full real-time reconstruction in Run 2 of the LHC*, JINST **14** (2019) P04013, arXiv:1812.10790.
- [33] LHCb collaboration, *LHCb RICH: Technical Design Report*, CERN-LHCC-2000-037, 2000.
- [34] J. L. Borges, *A universal history of infamy*, Penguin Books, Harlow, England, 1975.
- [35] A. Keune and G. Conti, *Magnetic field parametrisations*, CERN-LHCb-INT-2010-002, CERN, Geneva, 2010.
- [36] ROOT, *TMultiDimFit Class Reference*, 2024. root.cern.ch/root/html/TMultiDimFit.html.
- [37] J. Prisciandaro, F. Blanc, and T. Nakada, *Improved magnetic field map with 2011 measurements for the LHCb dipole magnet*, CERN-LHCb-INT-2012-012, CERN, Geneva, 2012.
- [38] M. A. Schubiger *et al.*, *Magnetic field map with 2014 measurements for the LHCb dipole magnet*, CERN-LHCb-INT-2015-034, CERN, Geneva, 2015.
- [39] P. Sainvitu and A. Zemanek, *Fiducialisation of the “sensors bar of the magnetic field bench” in the building 164*, EDMS document No. 2470215, CERN, 2021.
- [40] Dassault Systèmes, *Opera*, 2019.
<https://www.3ds.com/products-services/simulia/products/oper/>.
- [41] Dassault Systèmes, *Opera-3d reference manual*, version 2020.
- [42] P. Sainvitu and A. Zemanek, *Measurement of the ‘dipole inner planes’*, EDMS document No. 2595721, CERN, 2021.
- [43] M. Boruchowski and D. Mergelkuhl, *Photogrammetric measurements of relative dipole’s deformations, caused by different magnetic field conditions*, EDMS document No. 1231521, CERN, 2012.
- [44] P. Billoir, *Correcting a magnetic field map through the alignment of tracks*, Nucl. Instrum. Meth. **A902** (2018) 33.
- [45] Particle Data Group, S. Navas *et al.*, *Review of particle physics*, Phys. Rev. **D110** (2024) 030001.
- [46] R. Kopecna, *Tracking and vertexing in LHCb*, PoS **VERTEX2018** (2019) 039.

Bibliography

- [47] M. Calvo Gomez *et al.*, *A tool for γ/π^0 separation at high energies*, LHCb-PUB-2015-016, 2015.
- [48] D. Červenkov, *PIDCalib2*, 2022. <https://pypi.org/project/pidcalib2>.
- [49] LHCb collaboration, *GammaPi0SeparationCalib*, 2024. <https://gitlab.cern.ch/LHCb-CalObjects/gammapi0separationcalib>.
- [50] M. Kucharczyk, P. Morawski, and M. Witek, *Primary vertex reconstruction at LHCb*, LHCb-PUB-2014-044, 2014.
- [51] T. Fawcett, *An introduction to ROC analysis*, Pattern Recognit. Lett. **27** (2006) 861.
- [52] T. Skwarnicki, *A study of the radiative cascade transitions between the Upsilon-prime and Upsilon resonances*, PhD thesis, Institute of Nuclear Physics, Krakow, 1986, DESY-F31-86-02.
- [53] Particle Data Group, M. Tanabashi *et al.*, *Review of particle physics*, Phys. Rev. **D98** (2018) 030001.
- [54] ARGUS collaboration, H. Albrecht *et al.*, *Search for hadronic $b \rightarrow u$ decays*, Phys. Lett. **B241** (1990) 278.
- [55] BESIII collaboration, M. Ablikim *et al.*, *Coupled-channel analysis of the $\chi_{c1}(3872)$ line shape with BESIII data*, Phys. Rev. Lett. **132** (2024) 151903.
- [56] CMS collaboration, A. M. Sirunyan *et al.*, *Measurement of $B_c(2S)^+$ and $B_c^*(2S)^+$ cross section ratios in proton-proton collisions at $\sqrt{s} = 13$ TeV*, Phys. Rev. **D102** (2020) 092007.
- [57] LHCb collaboration, B. Adeva *et al.*, *Roadmap for selected key measurements of LHCb*, LHCb-PUB-2009-029, 2010.
- [58] M. Pivk and F. R. Le Diberder, *sPlot: A statistical tool to unfold data distributions*, Nucl. Instrum. Meth. **A555** (2005) 356, [arXiv:physics/0402083](https://arxiv.org/abs/physics/0402083).
- [59] H. Dembinski, M. Kenzie, C. Langenbruch, and M. Schmelling, *Custom orthogonal weight functions (COWs) for event classification*, Nucl. Instrum. Meth. **A1040** (2022) 167270, [arXiv:2112.04574](https://arxiv.org/abs/2112.04574).
- [60] F. James, *Statistical methods in experimental physics*, World Scientific, 2nd ed., 2006.
- [61] G. Breit and E. Wigner, *Capture of slow neutrons*, Phys. Rev. **49** (1936) 519.
- [62] G. J. Gounaris and J. J. Sakurai, *Finite-width corrections to the vector-meson-dominance prediction for $\rho \rightarrow e^+ e^-$* , Phys. Rev. Lett. **21** (1968) 244.
- [63] C. Zemach, *Use of angular-momentum tensors*, Phys. Rev. **140** (1965) B97.
- [64] S. U. Chung, *General formulation of covariant helicity-coupling amplitudes*, Phys. Rev. **D57** (1998) 431.

[65] T. Evans, B. Meadows, and G. Wilkinson, *Amplitude study of the decays $D \rightarrow K\pi\pi\pi$* , LHCb-ANA-2016-043 (2016).

[66] T. Evans, *Studies of the decays $D^0 \rightarrow K^\mp\pi^\pm\pi^\pm\pi^\mp$ at CLEO-c and LHCb*, PhD thesis, Oxford U., 2017.

[67] *AmpGen*, 2022. <https://github.com/GooFit/AmpGen>.

[68] W. D. Hulsbergen, *Decay chain fitting with a Kalman filter*, Nucl. Instrum. Meth. **A552** (2005) 566, arXiv:physics/0503191.

[69] V. Bellée, *Amplitude analysis for the measurement of the photon polarisation in $B \rightarrow K\pi\pi\gamma$ decays*, PhD thesis # 10183, EPFL, 2020.

[70] LHCb collaboration, R. Aaij *et al.*, *Amplitude analysis of $B^+ \rightarrow \psi(2S)K^+\pi^+\pi^-$ decays*, JHEP **01** (2025) 054, arXiv:2407.12475.

[71] LHCb collaboration, R. Aaij *et al.*, *Search for CP violation through an amplitude analysis of $D^0 \rightarrow K^+K^-\pi^+\pi^-$ decays*, JHEP **02** (2019) 126, arXiv:1811.08304.

[72] LHCb collaboration, R. Aaij *et al.*, *Studies of the resonance structure in $D^0 \rightarrow K^\mp\pi^\pm\pi^+\pi^-$ decays*, Eur. Phys. J. **C78** (2018) 443, arXiv:1712.08609.

[73] LHCb collaboration, R. Aaij *et al.*, *Measurement of the CKM angle γ and B_s^0 - \bar{B}_s^0 mixing frequency with $B_s^0 \rightarrow D_s^\mp h^\pm\pi^\pm\pi^\mp$ decays*, JHEP **03** (2021) 137, arXiv:2011.12041.

[74] M. Borsato *et al.*, *Calibration samples of γ and π^0 for particle identification at LHCb in 2017 and 2018*, CERN-LHCb-INT-2021-002, CERN, Geneva, 2021.

[75] Belle collaboration, H. Guler *et al.*, *Study of the $K^+\pi^-\pi^+$ final state in $B^+ \rightarrow J/\psi K^+\pi^+\pi^-$ and $B^+ \rightarrow \psi' K^+\pi^+\pi^-$* , Phys. Rev. **D83** (2011) 032005.

[76] CLEO collaboration, D. M. Asner *et al.*, *Hadronic structure in the decay $\tau^- \rightarrow \nu_\tau\pi^-\pi^0\pi^0$ and the sign of the tau neutrino helicity*, Phys. Rev. **D61** (1999) 012002.

[77] CMD-2 collaboration, R. R. Akhmetshin *et al.*, *Measurement of $e^+e^- \rightarrow \pi^+\pi^-$ cross-section with CMD-2 around ρ -meson*, Phys. Lett. **B527** (2002) 161, arXiv:hep-ex/0112031.

[78] C. Hanhart, *Amplitude analysis for mesons and baryons: Tools and technology*, AIP Conf. Proc. **1735** (2016) 020015.

[79] P. D'Argent, E. M. Gersabeck, and M. Pappagallo, *Amplitude analysis of $B^+ \rightarrow \psi(2S)K^+\pi^+\pi^-$ decays*, LHCb-ANA-2023-048, 2023.

[80] LASS collaboration, *A study of $K^-\pi^+$ scattering in the reaction $K^- p \rightarrow K^-\pi^+ n$ at 11 GeV/c*, Nucl. Phys. **B296** (1988) 493.

[81] W. Dunwoodie, *Fits to $K\pi$ $I = \frac{1}{2}$ S-wave amplitude and phase data*, 2013. http://www.slac.stanford.edu/wmd/kpi_swave/kpi_swave_fit.note.

Bibliography

- [82] S. M. Roy, *Exact integral equation for pion-pion scattering involving only physical region partial waves*, Phys. Lett. **B36** (1971) 353.
- [83] D. V. Bugg, *The mass of the σ pole*, J. Phys. **G34** (2006) 151.

Marie Bachmayer

Particle physicist looking for fresh challenges.

I want to work with motivated people towards a sustainable future.

mariebachmayer@gmail.com

+41 76 505 12 95

Engelstrasse 60, 8004 Zürich



LEBENSLAUF

PhD Physics

EPFL, CERN

September 2019 - February 2025

- **Data analysis** at the LHCb experiment at CERN, using machine learning techniques for the precise measurement of a Standard Model parameter
- Development of a novel magnetic field map of the LHCb dipole in exchange with the >1000 member strong **LHCb collaboration**
- **Science communication** at EPFL and at CERN: detector cavern tours, masterclasses, public events
- **Head teaching assistant** in various physics lectures: preparation of exercises, group tutoring, exam correction

MSc High Energy Physics

ETH Zürich - Université Paris-Saclay
with distinction

September 2017 - September 2019

GPA: 5.79

- **Joint Master** with one year each at ETH Zürich and École Polytechnique (Université Paris-Saclay, now Institut polytechnique de Paris)
- Tutoring: Mathematical Methods for Physics

Plaza Klub Zürich

Coat check

November 2015 - June 2016

BSc Physics

ETH Zürich

September 2012 - September 2017

GPA: 5.31

- Exchange semester at the **National University of Singapore**, August - December 2016
- **Teaching assistant** Linear Algebra: Mini-lectures for groups of approx. 30 students, feedback on assignments
- **VMP Assistentenaward** 2017: awarded prize for best teaching assistant by the student association
- Preparation and teaching of two one-week **exam preparation courses**

LANGUAGES

German (native)

English (fluent)

French (fluent)

PROGRAMMING LANGUAGES

Python (Data Science): proficient

C++ : basic