

Real-time analysis for studies of charge parity and lepton flavour violation in *b*-meson decays at LHCb

James Andrew Gooding
geboren in Leicester, Vereinigtes Königreich

Dissertation zur Erlangung des akademischen Grades
Dr. rer. nat.

Fakultät Physik
Technische Universität Dortmund

Dortmund, March 2026

Der Fakultät Physik der Technischen Universität Dortmund zur Erlangung des akademischen Grades eines Doktors der Naturwissenschaften vorgelegte Dissertation.

Gutachter:

Prof. Dr. Johannes Albrecht

Prof. Dr. Kevin Kröninger

Vorsitzender der Prüfungskommission:

Prof. Dr. Zhe Wang

Vertreter der wissenschaftlichen Mitarbeiter:

JProf. Dr. Benedikt Fauseweh

Datum der Einreichung der Dissertation: 5. September 2025

Datum der mündlichen Prüfung: 10. November 2025

This work is part of SMARTHEP, which received funding from the European Union's Horizon 2020 research and innovation programme under Grant Agreement n. 956086.

Abstract

Trigger systems of high energy physics experiments are crucial to the processing of the immense quantities of data produced from particle collisions. The LHCb experiment underwent a major upgrade between Runs 2 and 3 of the LHC, in which the trigger system was redeveloped from a hardware-and software-based model to an entirely software-based model. This thesis presents measurements with proton-proton collision data recorded by the LHCb experiment for both of these models, and the commissioning of the software-based trigger.

The first ever measurement of the time-integrated untagged CP asymmetry in the decay of $B_s^0 \rightarrow D_s^- \pi^+$, $\langle A_{\text{untagged}}^s \rangle$, is presented, performed with 5.4 fb^{-1} of data recorded in 2016-18. The measured values in the $D_s^- \rightarrow h^- h^+ \pi^-$ decay modes of the D_s^- , are consistent with the Standard Model of particle physics, down to a precision of $\mathcal{O}(10^{-3})$. An ongoing measurement of the lepton flavour universality ratio in $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ decays, $r_{J/\psi}$, using 4.5 fb^{-1} of data recorded in 2024 is then presented. This measurement validates the techniques used to evaluate efficiencies and data-simulation differences on these newest datasets, ahead of analyses in the rare $B^+ \rightarrow K^+ \ell^+ \ell^-$ decay modes. The impact of removing the hardware trigger is discussed throughout the thesis; the realisation of the anticipated improvements in trigger efficiencies is demonstrated explicitly.

Kurzfassung

Die Triggersysteme von Hochenergiephysikexperimenten sind entscheidend für die Verarbeitung der immensen Datenmengen, die bei Teilchenkollisionen anfallen. Das LHCb-Experiment wurde zwischen den Runs 2 und 3 des LHC eines umfassenden Upgrades unterzogen, bei dem das Triggersystem von einem hardware- und softwarebasierten Modell zu einem vollständig softwarebasierten Modell weiterentwickelt wurde. In dieser Arbeit werden Messungen mit Daten aus Proton-Proton-Kollisionen vorgestellt, die vom LHCb-Experiment für beide Modelle aufgezeichnet wurden, sowie die Inbetriebnahme des softwarebasierten Triggers.

Die allererste Messung der zeitintegrierten ungetaggtten CP -Asymmetrie im Zerfall von $B_s^0 \rightarrow D_s^- \pi^+$, $\langle A_{\text{untagged}}^s \rangle$, wird vorgestellt, die mit 5.4 fb^{-1} an Daten aus den Jahren 2016-18 durchgeführt wurde. Die gemessenen Werte in den $D_s^- \rightarrow h^- h^+ \pi^-$ -Zerfallsmoden des D_s^- sind mit dem Standardmodell der Teilchenphysik konsistent, und mit einer Genauigkeit von $\mathcal{O}(10^{-3})$. Anschließend wird eine laufende Messung des Leptonenflavour-Universalitätsverhältnisses in $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ -Zerfällen, $r_{J/\psi}$, unter Verwendung von 4.5 fb^{-1} an Daten aus dem Jahr 2024 vorgestellt. Diese Messung validiert die Techniken, die zur Bestimmung der Effizienz und der Unterschiede zwischen Daten und Simulationen bei diesen neuesten Datensätzen verwendet werden, bevor die Analysen der seltenen $B^+ \rightarrow K^+ \ell^+ \ell^-$ -Zerfälle durchgeführt werden. Die Auswirkungen der Abschaffung des Hardware-Triggers werden in der gesamten Arbeit erörtert; die Realisierung der erwarteten Verbesserungen der Trigger-Effizienz wird ausdrücklich nachgewiesen.

Acknowledgements

The last three years have been the most involved and eventful, yet exciting and memorable of my life. This would not have been possible without all the support I have received at home, at work and from friends around the world. Were the full list of those deserving thanks to be included here, the thesis would quite easily double in length—it truly takes a village. I would therefore like to dedicate this thesis to everybody who has helped me along the way. Sincerely, thank you.

I would also like to extend a few specific thanks. First and foremost, I would like to thank Johannes for the supervision, guidance and advice over these three years, which have undoubtedly made me a better physicist than the fresh-faced student who stepped off of the plane in the summer of 2022.

Likewise, Ale and Biljana, your support and guidance throughout the past years have truly helped me to overcome the hurdles of those years' work. Ale, I can only apologise for the many draft thesis chapters you have read and re-read, I cannot thank you enough for enduring my writing! I'd also like to extend my thanks (and indeed apologies) to Quentin, Nicole (both of you!), Vukan, James and Noah, who very generously gave their time to help this thesis take its form.

I am extremely grateful to have been able to work on the analyses and projects discussed in this thesis with such excellent colleagues and friends. It has been a pleasure to work with all of you thus far, and I look forward to working with you all in the future.

I owe a great deal of thanks to many people whose advice I have asked over the past years. Nicole, Agnieszka, Vava and Conor, it would be amiss if I did not thank you here; your advice has made sure that I take pride in everything I do, has kept me motivated to work towards my goals, and made me excited to see what the future brings.

Furthermore, I have been blessed with fantastic family and friends; in Dortmund, in the UK, and across the world. Your humour and kindness have kept me sane and brought a smile to my face countless times.

Finally, and arguably most importantly, I would like to thank those closest to me for their support through the many highs and lows of these challenging but ultimately extremely rewarding times. To Vanessa, to my family, with all my heart, thank you. Without your care, patience and interest I would not have been able to strive for my goals and dreams; there are no words with which I can express how grateful I am.

Contents

1	Introduction	1
2	Heavy flavour physics in the Standard Model	3
2.1	Fundamental particles of the SM	3
2.2	Symmetries and interactions of the SM	4
2.2.1	Symmetries, local gauge invariance and QCD	4
2.2.2	Electroweak unification and spontaneous symmetry breaking	7
2.2.3	Weak currents, the CKM matrix and flavour-changing processes	9
2.2.4	Phenomenology of heavy flavour physics	11
2.2.5	Effective field theories for heavy flavour phenomenology	16
2.3	Looking beyond the SM	17
2.4	Current experimental status	19
2.4.1	The $b \rightarrow c\bar{u}q$ anomaly	19
2.4.2	Lepton flavour anomalies	20
3	The LHCb experiment during Run 2 of the LHC	21
3.1	The Large Hadron Collider	21
3.2	The LHCb detector in Run 2	22
3.2.1	Tracking of charged particles and reconstruction of vertices	23
3.2.2	Particle identification and calorimetry	26
3.2.3	The LHCb data flow	28
4	Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays	35
4.1	Analysis strategy, samples and selections	35
4.1.1	Data samples and trigger selections	37
4.1.2	Background contributions	37
4.1.3	Offline selection	40
4.1.4	Studies of possible remaining backgrounds	44
4.1.5	Replication of selection on MC samples	45
4.2	Estimation of efficiencies	45
4.2.1	Kinematic selection efficiencies	47
4.2.2	PID efficiencies	48
4.2.3	L0 trigger efficiencies	49
4.2.4	Single-pion tracking efficiencies	50
4.2.5	Efficiencies as per-event weights	51
4.3	Raw asymmetry, A_{raw}	52
4.3.1	Modelling of signal and background components	53
4.3.2	The nominal invariant mass fit	58
4.3.3	Fit validation and pseudoexperiments	60
4.3.4	Asymmetry contributions from background components	63
4.4	Detection asymmetry, $A_{\text{det.}}$	67
4.4.1	Contributions from kinematic selection efficiencies	68
4.4.2	Contribution from PID efficiencies	68
4.4.3	Contribution from L0 trigger efficiencies	69
4.4.4	Tracking asymmetry	70
4.4.5	Summary of $A_{\text{det.}}$	73

4.5	Production asymmetry, A'_{prod}	73
4.6	Systematic uncertainties	73
4.6.1	Systematic effects in A_{raw}	74
4.6.2	Systematic effects in A_{det}	75
4.7	Consistency of $\langle A_{\text{untagged}}^s \rangle$ in data subsamples	77
4.8	Summary of results and conclusion	77
5	A 30 MHz LHCb software trigger for Run 3	79
5.1	The LHCb experiment during Run 3 of the LHC	79
5.1.1	Upgrades to the LHCb detector	79
5.1.2	Upgrade of the LHCb trigger	81
5.1.3	Status of Run 3 data-taking	88
5.2	Development and performance of selections for the Run 3 software trigger	89
5.2.1	Development of inclusive triggers for HLT2	89
5.2.2	Estimating trigger efficiencies	102
5.2.3	Inclusive trigger efficiencies in 2024 and 2025 data-taking	110
5.3	Conclusion	116
6	Measurement of $r_{J/\psi}$ in 2024 LHCb data	117
6.1	Analysis strategy, samples and selections	117
6.1.1	Data samples and trigger selections	118
6.1.2	Background contributions	119
6.1.3	Offline selection	119
6.2	Extraction of yields	121
6.3	Estimation of efficiencies and data-MC corrections	123
6.3.1	Geometric acceptance efficiencies	125
6.3.2	Data-MC corrections	125
6.3.3	Combination of efficiencies and corrections	130
6.4	Consideration of systematic uncertainties	132
6.5	Summary of results and conclusion	134
7	Conclusion	135
	Appendices	138
A	Appendix for the Run 2 measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$	139
A.1	Updated results	139
A.2	Per-event weighting of MC samples	139
A.3	Parameterisation and results of the invariant mass fits	142
A.3.1	Signal models	142
A.3.2	Combinatorial models	145
A.3.3	Results of the nominal invariant mass fit	145
B	Appendix to the LHCb trigger in Run 3	147
B.1	Selections of the cut-based inclusive detached dilepton trigger	147
B.2	Retuning of the topological b -trigger	150
C	Appendix to the 2024 measurement of $r_{J/\psi}$	153
C.1	Requirements of exclusive HLT2 Turbo lines	153

C.2	Extraction of yields	153
C.2.1	Dielectron mass resolution smearing	156
	List of Tables	156
	List of Figures	158
	Bibliography	163

Two roads diverged in a yellow wood,
And sorry I could not travel both
And be one traveler, long I stood
And looked down one as far as I could
To where it bent in the undergrowth;

Then took the other, as just as fair,
And having perhaps the better claim,
Because it was grassy and wanted wear;
Though as for that the passing there
Had worn them really about the same,

And both that morning equally lay
In leaves no step had trodden black.
Oh, I kept the first for another day!
Yet knowing how way leads on to way,
I doubted if I should ever come back.

I shall be telling this with a sigh
Somewhere ages and ages hence:
Two roads diverged in a wood, and I—
I took the one less traveled by,
And that has made all the difference.

“The Road Not Taken”, *Robert Frost*

1 Introduction

The Standard Model (SM) of particle physics is a relativistic quantum field theory, which has thus far provided accurate and robust descriptions of the elementary particles of our universe and their interactions. Whilst broadly successful, this theory, introduced in Chapter 2, does not describe all known physical phenomena. For example, gravitational interactions are not described in the SM, and its theory does not predict suitable candidates for dark matter, the presence of which is supported by extensive cosmological evidence. It is on this basis that the focus of modern particle physics turns to the hunt for new physics (NP), that is, physics beyond the SM.

In this field, the decays of mesons containing a heavy b -quark are of particular interest. Their study can simultaneously probe the properties of the weak interaction of the SM and seek for possible effects from NP. These decays often have clear experimental signatures from which they can be detected and clean theoretical predictions which leverage the large mass of the b -quark. In combination, these allow for powerful comparisons of measured observables against their corresponding SM predictions.

The Large Hadron Collider beauty (LHCb) experiment at the Large Hadron Collider (LHC), described in Chapter 3, was constructed with the purpose of studying such properties and decays. To this end, the LHCb detector was designed around the efficient reconstruction of trajectories of charged particles, their high-quality identification and the precise determination of the positions of their corresponding vertices. This is achieved through a unique detector geometry, lying entirely in the forward direction (in contrast to the hermetic designs typical of general purpose detectors), a large dipole magnet to deflect charged particles, and dedicated tracking and particle identification subsystems. Additionally, the large luminosities achieved by the high energy proton-proton (pp) collisions at the LHC provide the largest sample of b -hadron decays in the world. These make LHCb an ideal observatory to seek NP in b -hadron decays.

Not every detected product of every pp collision will be of interest to the LHCb physics programme; of the visible interactions occurring at 30 MHz, only a small fraction can be feasibly recorded. To handle the vast quantity of information produced by the detector, LHCb employs a trigger system to process and filter the data. In Run 2 of the LHC, this was a hardware- and software-based system, described in Chapter 3; for Run 3 of the LHC, the LHCb detector was upgraded and an entirely software-based trigger system developed, described in Chapter 5.

This thesis presents two measurements performed in LHCb pp collision data. The first, described in Chapter 4, is a measurement of CP asymmetry in $B_s^0 \rightarrow D_s^- \pi^+$ decays in data collected with 2016-18, *i.e.*, with the hardware- and software-based trigger system. The second, described in Chapter 6, is a measurement of lepton flavour universality (LFU) in $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ decays with data collected in 2024, *i.e.*, with the upgraded trigger system. The upgrade of the LHCb detector and trigger, which took place between the collection of these datasets is described at length in Chapter 5. Both the measurements and upgrades were carried out by groups of researchers from the LHCb collaboration, with the author of this thesis having made significant contributions. It is on these contributions which this thesis focuses.

In Chapter 4, the measurement of the time-integrated untagged CP asymmetry in $B_s^0 \rightarrow D_s^- \pi^+$ decays, on which the author of this thesis is a proponent and contact

1 Introduction

author, is presented. This measurement was in the late stages of internal LHCb review at the time of submission the thesis and is presented accordingly. This measurement has since been submitted to *Physical Review Letters*, with the updated result briefly presented in Appendix A.1 Tensions between experimental measurements of the branching fractions of $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays and their corresponding theoretical predictions have emerged in recent years; contributions from NP effects are among the possible explanations offered for these tensions, as discussed in Chapter 2. Such a contribution could induce CP violation in the decays of the tree level $B^0 \rightarrow D^- K^+$ and $B_s^0 \rightarrow D_s^- \pi^+$ decays; the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays, $\langle A_{\text{untagged}}^s \rangle$, is thus measured in 2016-18 pp collision data. This is the first measurement of $\langle A_{\text{untagged}}^s \rangle$, providing unique constraints on potential NP contributions to $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays. The measurement is statistically limited, with the L0 hardware trigger causing the greatest reduction in statistics; a significant improvement in the sensitivity of LHCb to CP violation from such a NP contribution could be expected from the removal of L0 in the LHCb upgrade.

The software-only trigger system is the focus of Chapter 5, which describes the upgrade of the LHCb detector and discusses the development and evaluation of selection algorithms in the upgraded trigger. In this chapter, the processes of developing selection algorithms in HLT2 are laid out, illustrated by the examples of the cut-based inclusive detached dilepton trigger and the retuning of the topological b -hadron trigger. A data-driven approach for evaluating trigger efficiencies is then presented; in particular, the implementation of this approach in the `TriggerCalib` common software tool (which the author of this thesis is the primary author of) is discussed. This approach is applied to evaluate the efficiencies of inclusive HLT1 and HLT2 selections in pp collision data collected in 2024 and 2025, which are compared against the corresponding efficiencies in Run 2 (incorporating the efficiency of the L0 trigger).

Finally, Chapter 6 presents a measurement of the lepton flavour universality ratio of $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ decays, $r_{J/\psi}$, in 2024 pp collision data. This measurement is ongoing and is reported per its status at the time of submission of the thesis. The observable $r_{J/\psi}$ is a precursor to the equivalent observable in the rare $B^+ \rightarrow K^+ \ell^+ \ell^-$ modes, R_K , and provides an important validation to complex analyses in the LFU sector. Furthermore, $r_{J/\psi}$ can be incorporated directly, writing R_K as a double ratio to significantly reduce the effect of systematic uncertainties. As $r_{J/\psi}$ is predicted to be 1 to a high precision, any observed deviations from 1 can be concluded to arise from the analysis techniques applied to measure $r_{J/\psi}$. This makes $r_{J/\psi}$ an ideal observable in which to understand the features of the 2024 dataset, *i.e.*, those arising from the upgrade of the detector and trigger, ahead of analyses of LFU ratios in rare decays.

2 Heavy flavour physics in the Standard Model

The SM is a relativistic quantum field theory, providing the current best theoretical understanding of elementary particles and the interactions between them [1, 2, 3]. The SM, which describes three of the four fundamental forces, emerged from the great breakthroughs in quantum mechanics in the early 20th century and the later applications of group theory to quantum mechanics in the 1950s. The phenomena predicted by the SM have been consistently supported by experimental observations, with all of the particles predicted by the SM having now been discovered. Nonetheless, several questions about the nature of our universe and its contents remain unanswered by the SM, making its precise predictions a prime target for deeper interrogation.

This chapter lays out the fundamentals of the SM, demonstrates the successes and shortcomings of the SM, and the subsequent need for NP. In particular, this chapter highlights the decays of b -hadrons as a theoretically and experimentally favourable observatory for searches for new physics, which run orthogonal to direct searches.

2.1 Fundamental particles of the SM

The SM can be described mathematically by its gauge (local symmetry) group, the product of three constituent gauge groups,

$$\underbrace{\text{SU}(3)_C}_{\text{QCD}} \times \underbrace{\text{SU}(2)_L \times \text{U}(1)_Y}_{\text{Weak and EM}}. \quad (2.1)$$

The first term is the $\text{SU}(3)_C$ symmetry of the strong interaction or quantum chromodynamics (QCD), discussed in Section 2.2.1. The second and third terms describe the $\text{SU}(2)_L$ and $\text{U}(1)_Y$ symmetries of the electroweak (EW) theory, from which the weak and electromagnetic (EM) interactions arise, discussed in Section 2.2.2.

Each interaction is mediated by gauge bosons, which couple to specific properties of fundamental particles: the gluon, g , mediates the strong interaction and couples to colour charge; the W and Z bosons mediate the weak interaction, coupling to the weak isospin and hypercharge; the photon mediates the EM interaction, γ , couples to the electric charge. As bosons, each gauge boson carries an integer spin; all four mediators carry a spin of 1.

The matter of the universe is however not formed of these bosons, and instead is built from fermions: elementary particles carrying a half-integer spin. Fermions are classed as left- or right-handed based on whether their momenta and spin are in the opposite or same direction, respectively. The fermions can be split based on whether they carry a colour-charge: colour-charged fermions are quarks, colour-neutral fermions are leptons. The quarks can be subdivided further into the up (u), charm (c), and top (t) quarks (grouped together as up-type quarks), each of which have an electric charge (in units of the elementary charge, e) of $+2/3$, and the down (d), strange (s), and bottom (b) quarks (grouped together as down-type quarks), which have an electric charge of $-1/3$. The leptons can be similarly subdivided into the electrically charged leptons (with a charge of -1): the electron (e), muon (μ), and tau (τ); and their electrically neutral

counterparts, the neutrinos: the electron neutrino (ν_e), the muon neutrino (ν_μ), and the tau neutrino (ν_τ). The charged leptons and their corresponding neutrinos each share a conserved quantum number, the electron, muon and tau numbers, to which each ℓ^-/ν_ℓ contributes +1 and each $\ell^+/\bar{\nu}_\ell$ contributes -1. Each of the triplets of particles listed above are given in increasing mass order, except for the neutrinos, which are considered massless in the SM and to which the mass ordering is currently an open question [4]

The final fundamental particle of the SM to be discovered, which completes the SM, is the Higgs boson: a scalar (carrying a spin of 0) boson, with no colour or electric charge but a large mass of (125.20 ± 0.11) GeV [5], which couples to massive particles but does not convey a fundamental force. The coupling of the Higgs boson to massive particles generates their masses (see Section 2.2.2). The Higgs boson was discovered jointly by the ATLAS and CMS experiments only as recently as 2012 [6, 7], almost 50 years after having been postulated [8, 9, 10, 11].

The full set of particles of the SM are shown (excluding antiparticles) in Fig. 2.1. In the next sections, we consider in more detail the interactions between these particles.

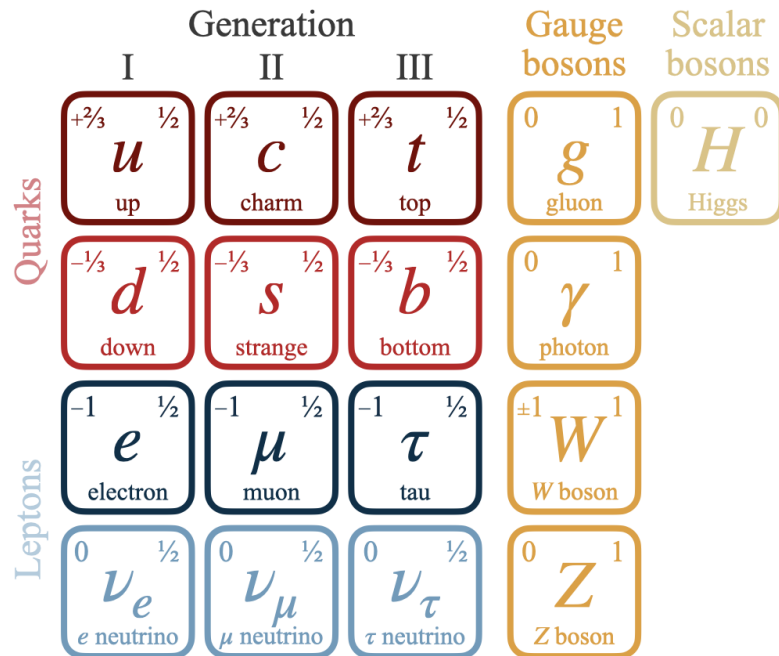


Figure 2.1: The fundamental particles of the SM, with electric charges (upper left) and spins (upper right) annotated.

2.2 Symmetries and interactions of the SM

The SM is defined by the symmetries of its quantum fields, which generate the interactions between particles. This section introduces how each of the interactions arises and lays out the ensuing properties and behaviours of b -hadrons under these interactions.

2.2.1 Symmetries, local gauge invariance and QCD

Symmetries under discrete and continuous local (gauge) transformations are the cornerstones of the SM. The former define how interactions may proceed, whilst the latter generate the interactions themselves.

The SM is subject to three discrete symmetries, each of which manipulate a state $\psi(\vec{x}, t)$. Conjugation of charges, C , inverts all charges, *i.e.*, $Q \xrightarrow{C} -Q$. Parity inversion, P , reverses all spatial directions, *i.e.*, $\psi(\vec{x}, t) \xrightarrow{P} \psi(-\vec{x}, t)$. Time reversal, T , reverses the direction of time $\psi(\vec{x}, t) \xrightarrow{T} \psi(\vec{x}, -t)$. These symmetries can be combined, for example, the CP symmetry which transforms particles to antiparticles and vice versa. The combination of CP and T symmetries must not be violated under the CPT theorem [12].

The EM interaction conserves each of the C , P and T symmetries, and their combinations. The same is not strictly true for the strong and weak interactions, though violation of these symmetries by the strong interaction has not been observed. In 1956, the Wu experiment, examining the parity of electrons produced in the β^- decay of ^{60}Co , demonstrated that, not only does the weak interaction violate P symmetry, but that it does so maximally [13]. This was followed in 1964 by evidence of CP violation in the weak decays of neutral kaons [14]. The K^0 is observed in short- and long-lived states (see Section 2.2.4), which decay predominantly to $\pi^{0(+)}\pi^{0(-)}$ (CP even) and $\pi^{0(+)}\pi^{0(-)}\pi^0$ (CP odd) final states, respectively. On this basis, these were thought to be CP eigenstates; however, the decay of the long-lived state to $\pi^0\pi^0$ was observed, which is only possible if the long-lived state is not a CP eigenstate, *i.e.*, if the CP symmetry is violated by the weak interaction.

According to Noether's theorem, every continuous symmetry has a corresponding conservation law [15]. The local gauge transformations of the SM, those described in Eq. 2.1, correspond to the invariance of the Lagrangian density under those transformations. It is from this property that the interactions of the SM arise.

To understand how these interactions arise, we consider a free massive fermion, the kinematics of which are described, according to the Dirac equation, by a Lagrangian density of

$$\mathcal{L} = \bar{\phi}(i\gamma^\mu\partial_\mu - m)\phi. \quad (2.2)$$

As an example, we can enforce that this has local gauge invariance under a $U(1)$ symmetry, *i.e.*, under a simple local rotation of the field, $\phi \rightarrow e^{iq\alpha(x)}\phi$. Applying this transformation, it quickly becomes apparent that Eq. 2.2 is not invariant under $U(1)$. To rectify this, a covariant derivative is defined as

$$D_\mu = \partial_\mu + iq\alpha(x)A_\mu(x), \quad (2.3)$$

where q describes the scale of transformation (later the coupling strength) and $A_\mu(x)$ is a gauge field corresponding to the generator of $U(1)$, which transforms as

$$A_\mu \rightarrow A_\mu - \frac{1}{q}\partial_\mu\alpha. \quad (2.4)$$

Replacing the partial derivative in Eq. 2.2 with this covariant derivative, the Lagrangian density becomes

$$\mathcal{L} = \bar{\phi}(i\gamma^\mu D_\mu - m)\phi. \quad (2.5)$$

A corresponding field strength tensor, $F_{\mu\nu}$ can be written as

$$F_{\mu\nu} = \frac{1}{iq}[D_\mu, D_\nu] = \partial_\mu A_\nu - \partial_\nu A_\mu, \quad (2.6)$$

such that a term describing the kinematics of the gauge field can be added to the Lagrangian density:

$$\mathcal{L} = -\frac{1}{4}F_{\mu\nu}F^{\mu\nu} + \bar{\phi}(i\gamma^\mu D_\mu - m)\phi. \quad (2.7)$$

2 Heavy flavour physics in the Standard Model

We can apply this approach to QCD, which is described by the $SU(3)_C$ symmetry group, wherein the C signifies that the strong interaction couples to colour charges. The charge conjugates of these colours are the anticolours \bar{r} , \bar{g} , and \bar{b} , respectively; the combination of a like colour and anticolour, *e.g.*, $r\bar{r}$ is considered a colour-neutral state. Quarks carry a single colour charge, antiquarks carry a single anticolour, and gluons carry both a colour and anticolour charge. A $SU(N)$ symmetry group has $N^2 - 1$ generators, thus the $SU(3)_C$ symmetry group has 8 generators, the 8 Gell-Mann matrices, λ_a , which are used to describe the 8 independent gluon colour configurations, g^a , and, by extension, the gluon fields A_μ^a . The coupling strength in QCD is $g_s = \sqrt{4\pi\alpha_s}$. A covariant derivative is obtained as in the example above; however, with the additional aspect of specifying colours i, j :

$$(D_\mu)_{ij} = \partial_\mu \delta_{ij} - ig_s \frac{(\lambda_a)_{ij}}{2} A_\mu^a. \quad (2.8)$$

As in Eq. 2.6, the field strength tensor $G_{\mu\nu}^a$ can be expressed in terms of D_μ as

$$G_{\mu\nu}^a = \frac{1}{ig_s} [D_\mu, D_\nu] = \partial_\mu A_\nu^a - \partial_\nu A_\mu^a + g_s f^a_{bc} A_\mu^b A_\nu^c, \quad (2.9)$$

where f^a_{bc} are the structure constants of the $SU(3)$ group. From this field strength tensor and the covariant derivative, an analogous Lagrangian density can be written for the field of a quark of flavour f , q_f under QCD as

$$\mathcal{L}_{\text{QCD}} = -\frac{1}{4} G_{\mu\nu}^a G^{a,\mu\nu} + \bar{q}_f (i\gamma^\mu D_\mu - m) q_f. \quad (2.10)$$

The first term of the QCD Lagrangian density gives rise to the gluon self-coupling, containing terms describing the 3- and 4-gluon interactions. These self-interactions result in two phenomena, defining the behaviour of QCD at high and low energy scales μ (short and long distances), respectively. The strong coupling constant, $\alpha_s(\mu)$, describes the strength of the strong interaction at a given energy scale μ . As $\mu \rightarrow \infty$, $\alpha_s(\mu) \rightarrow 0$: at extremely short distances ($\lesssim 0.1$ fm), quarks and gluons act as if free, a phenomenon known as asymptotic freedom [16]. The energy scale at which asymptotic freedom is no longer applicable is Λ_{QCD} , which is typically taken to be ~ 200 MeV. Similarly, as $\mu \rightarrow 0$, *i.e.*, for a larger distance between two quarks, $\alpha_s(\mu) \rightarrow \infty$, which limits the strong interaction to a finite distance, that is, the separation at which the energy required to maintain the separation is sufficiently large to produce a $q\bar{q}$ pair, considered to be ~ 1 fm. This effect, known as colour confinement, prevents colour-charged states from being observed.

As quarks cannot be observed as free particles, composite particles known as hadrons, formed of combinations of quarks/antiquarks held together in colour-neutral bound states by the strong interaction, are observed instead.

The quarks which directly form a hadron are known as valence quarks, though additional gluons and $q\bar{q}$ pairs (sea quarks) appear transiently, arising from quantum fluctuations [17]. Hadrons are classified as either mesons or baryons, depending on the number of quarks contained: mesons contain the same number of quarks and antiquarks (as pair(s) of like colour and anticolour), and thus have integer spins; baryons contain an odd number of (anti)quarks and thus have half-integer spins. Typically, meson refers to $q\bar{q}$ states and baryon to $qqq/\bar{q}\bar{q}\bar{q}$ states; however, states containing 4- and 5- quarks/antiquarks have been observed, with the first pentaquark states observed at LHCb in 2015 [18].

2.2.2 Electroweak unification and spontaneous symmetry breaking

In the 1950s/60s, the EM and weak interactions were unified into a single theoretical framework, known as the EW interaction [1,2,3]. In this framework, the EW interaction is described by the product of two symmetry groups: $SU(2)_L \times U(1)_Y$, where L indicates that the $SU(2)$ group acts only on left-handed fermions and the Y corresponds to the weak hypercharge. The gauge fields $W_\mu^{1,2,3}$ and B_μ arise from the generators of the $SU(2)$ and $U(1)$ groups, respectively, with the corresponding covariant derivative

$$D_\mu = \partial_\mu - i\frac{g_w}{2}T^i W_\mu^i - i\frac{g'_w}{2}Y B_\mu, \quad (2.11)$$

where g_w and g'_w describe the strength of couplings to weak isospin and weak hypercharge, T^i are the components of the weak isospin (the Pauli matrices σ^i) and Y is the weak hypercharge, defined such that $Q = T^3 + Y/2$. The corresponding field strength tensors

$$W_{\mu\nu}^i = \partial_\mu W_\nu^i - \partial_\nu W_\mu^i + g_w \varepsilon^{ijk} W_\mu^j W_\nu^k, \quad (2.12a)$$

$$B_{\mu\nu} = \partial_\mu B_\nu - \partial_\nu B_\mu, \quad (2.12b)$$

contribute to the Lagrangian density as

$$\mathcal{L}_{\text{gauge}} = -\frac{1}{4}W_{\mu\nu}^i W^{i,\mu\nu} - \frac{1}{4}B_{\mu\nu} B^{\mu\nu}. \quad (2.13)$$

The gauge fields $W_\mu^{1,2,3}$ and B^μ are massless, as introducing a direct mass term for any one of these fields would violate local gauge invariance. However, we know experimentally that 3 of the 4 electroweak mediators (W^+ , W^- , Z and γ) are massive. The framework must reconcile this difference and produce the mediator masses by other means.

The solution to this, known as the Higgs-Englert-Brout (or Higgs) mechanism, was formulated in 1964 [8,9,10,11]. A complex scalar (Higgs) field, $\phi = (\phi^+, \phi^0)$ is introduced, with a corresponding Lagrangian of the form

$$\begin{aligned} \mathcal{L}_{\text{Higgs}} &= (D_\mu \phi)^\dagger (D^\mu \phi) - V(\phi) \\ &= \mu^2 \phi^\dagger \phi + \lambda (\phi^\dagger \phi)^2, \end{aligned} \quad (2.14)$$

where the first term is the kinematic term of the Higgs field and the second the Higgs potential, and wherein μ and λ are free complex parameters. For values of $\mu^2 < 0$, the potential is unstable at the origin and has infinitely many stable minima at $\phi^\dagger \phi = -\mu^2/2\lambda$.

The expectation of the Higgs field in the lowest energy state can therefore be written, $\phi = \frac{1}{\sqrt{2}}(0, v)$ where $v = \sqrt{-\mu^2/2\lambda}$ is known as the vacuum expectation value (VEV). As the Higgs field takes the VEV, the EW symmetry is spontaneously broken as a specific minimum of $\phi^\dagger \phi$ is chosen. In this case, the kinematic term becomes

$$(D_\mu \langle \phi \rangle)^\dagger (D^\mu \langle \phi \rangle) = \frac{v^2}{8} \left(g_w^2 (W_\mu^1)^2 + g_w^2 (W_\mu^2)^2 + (g_w W_\mu^3 - g'_w B_\mu)^2 \right). \quad (2.15)$$

To simplify this expression, four fields are defined as mixtures of W_μ^i and B_μ :

$$\begin{pmatrix} W_\mu^+ \\ W_\mu^- \end{pmatrix} = \frac{1}{\sqrt{2}} \begin{pmatrix} 1 & 1 \\ 1 & -1 \end{pmatrix} \begin{pmatrix} W_\mu^1 \\ W_\mu^2 \end{pmatrix}, \quad (2.16a)$$

2 Heavy flavour physics in the Standard Model

$$\begin{pmatrix} A_\mu \\ W_\mu \end{pmatrix} = \begin{pmatrix} \cos \theta_w & \sin \theta_w \\ -\sin \theta_w & \cos \theta_w \end{pmatrix} \begin{pmatrix} B_\mu \\ W_\mu^3 \end{pmatrix}, \quad (2.16b)$$

wherein θ_w is the weak mixing (Weinberg) angle, defined such that $\tan \theta_w = g'_w/g_w$ and which can be related to the e as $e = g_w \sin \theta_w = g'_w \cos \theta_w$.

The kinematic term can therefore be written in terms of these fields as

$$\begin{aligned} (D_\mu \langle \phi \rangle)^\dagger (D^\mu \langle \phi \rangle) &= \frac{g_w^2 v^2}{4} W_\mu^- W^{+\mu} + \frac{v^2 (g_w^2 + g_w'^2)}{8} Z_\mu Z^\mu \\ &= M_W^2 W_\mu^- W^{+\mu} + \frac{1}{2} M_Z^2 Z_\mu Z^\mu, \end{aligned}$$

where $M_W = g_w v/2$ and $M_Z = \sqrt{g_w^2 + g_w'^2} v/2$ are the masses of the W and Z bosons. These masses can be related via the Weinberg angle as $\cos \theta_w = M_W/M_Z$. No such mass term arises for A^μ , the gauge field of the γ , with the corresponding symmetry, $U(1)_{\text{EM}}$, remaining unbroken.

The EW interaction, by construction as a $SU(2)_L \times U(1)_Y$ gauge theory, treats left- and right-handed (anti)particles differently: $SU(2)_L$ acts only on left-handed fermions; the $U(1)_Y$ acts on both left- and right-handed fermions, according to their weak hypercharge. This treatment is described by representing the left-handed fermion fields as doublets, *i.e.*,

$$Q_L^i = \begin{pmatrix} u'_L \\ d'_L \end{pmatrix}, \begin{pmatrix} c'_L \\ s'_L \end{pmatrix}, \begin{pmatrix} t'_L \\ b'_L \end{pmatrix}, \quad (2.17a)$$

$$L_L^i = \begin{pmatrix} \nu'_{e,L} \\ e'_L \end{pmatrix}, \begin{pmatrix} \nu'_{\mu,L} \\ \mu'_L \end{pmatrix}, \begin{pmatrix} \nu'_{\tau,L} \\ \tau'_L \end{pmatrix}, \quad (2.17b)$$

for quarks and leptons, respectively, where the primes simply denote that these are flavour eigenstates of the weak interaction. The right-handed fermion fields are represented by singlets: $u_R^i = u'_R, c'_R, t'_R$ and $d_R^i = d'_R, s'_R, b'_R$ for quarks, and $e_R^i = e'_R, \mu'_R, \tau'_R$ for the charged leptons. No right-handed singlets are considered for neutrinos, as these are not part of the SM and remain as yet unobserved. The kinematic term of fermions in the EW Lagrangian density can be written in terms of these fields as

$$\mathcal{L}_{\text{kinematic}} = \sum_{q=Q_L^i, u_R^i, d_R^i} i \bar{q} \gamma^\mu D_\mu q + \sum_{\ell=L_L^i, e_R^i} i \bar{\ell} \gamma^\mu D_\mu \ell. \quad (2.18)$$

However, if fermions were to be inherently massive, much as for the mediators, the resulting mass terms would not be invariant under the EW gauge symmetry. We must therefore consider them to be massless and look again to the Higgs field to generate masses.

The couplings of the fermions to the Higgs field, known as the Yukawa interactions/couplings, prior to symmetry breaking, are given by

$$\mathcal{L}_{\text{Yukawa}} = -Y_u \bar{Q}_L^i \phi^* u_R^i - Y_d \bar{Q}_L^i \phi d_R^i - Y_e \bar{L}_L^i \phi^* e_R^i, \quad (2.19)$$

where Y_α are the 3×3 Yukawa matrices, describing the strength of each coupling, and where higher order terms are conventionally ignored. Upon symmetry breaking, these terms become conventional mass terms, *e.g.*, $\frac{v}{\sqrt{2}} Y_u \bar{u}_L u_R$, generating the fermion masses.

After symmetry breaking, the Yukawa matrices can be diagonalised by rotating from the flavour eigenstates, $u'_{L,R}, d'_{L,R}, e'_{L,R}$ to the corresponding mass eigenstates, $u_{L,R}, d_{L,R}$,

$e_{L,R}$. In the quark fields, these rotations are performed by the 3×3 unitary matrices $V_{L,R}^{u,d}$, e.g., $u'_L = V_L^u u_L$, though only one Yukawa matrix can be diagonalised at a time. If Y^d is diagonalised, then u'_L are considered equal to u_L , and the down-type flavour eigenstates d'_L are rotated to the equivalent mass eigenstates d_L , where the corresponding rotation matrix is in fact the Cabibbo-Kobayashi-Maskawa (CKM) matrix, discussed in Section 2.2.3.

2.2.3 Weak currents, the CKM matrix and flavour-changing processes

Flavour-changing processes can only proceed via the weak interaction. The weak interaction can be split into two classes of interactions: charged currents, mediated by the W^\pm bosons, and neutral currents, mediated by the Z boson. Charged currents involve both a change in electric charge of e and a change of flavour, between up-type and down-type quarks or between charged and neutral leptons. The contribution to the Lagrangian density of the charged current for quarks can be written as

$$\mathcal{L}_{\text{CC}}^q = \frac{g_w}{\sqrt{2}} \left(\bar{u}'_L \gamma_\mu W^{-\mu} d'_L + \bar{d}'_L \gamma_\mu W^{+\mu} u'_L \right), \quad (2.20)$$

in terms of the weak flavour eigenstates. Rotating from the weak flavour eigenstates to the mass eigenstates per Section 2.2.2, this can also be written as

$$\mathcal{L}_{\text{CC}}^q = \frac{g_w}{\sqrt{2}} \left(\bar{u}_L (V_L^u)^\dagger V_L^d \gamma_\mu W^{-\mu} d_L + \bar{d}_L (V_L^d)^\dagger V_L^u \gamma_\mu W^{+\mu} u_L \right), \quad (2.21)$$

wherein $(V_L^u)^\dagger V_L^d$ is the CKM matrix V_{CKM} , and likewise $(V_L^d)^\dagger V_L^u$ is V_{CKM}^\dagger . The CKM matrix thus determines the relative probabilities with which the flavour-changing transitions between up- and down-type quarks proceed, with the complex CKM elements labelled according to these transitions:

$$V_{\text{CKM}} = \begin{pmatrix} V_{ud} & V_{us} & V_{ub} \\ V_{cd} & V_{cs} & V_{cb} \\ V_{td} & V_{ts} & V_{tb} \end{pmatrix}. \quad (2.22)$$

The elements of this matrix can be parameterised in many ways, though commonly these are expressed in terms of three real mixing angles and a complex phase [19]:

$$\begin{aligned} V_{\text{CKM}} &= \begin{pmatrix} 1 & 0 & 0 \\ 0 & c_{23} & s_{23} \\ 0 & -s_{23} & c_{23} \end{pmatrix} \begin{pmatrix} c_{13} & 0 & s_{13} e^{-i\delta_{13}} \\ 0 & 1 & 0 \\ -s_{13} & 0 & c_{13} e^{i\delta_{13}} \end{pmatrix} \begin{pmatrix} c_{12} & s_{12} & 0 \\ -s_{12} & c_{12} & 0 \\ 0 & 0 & 1 \end{pmatrix} \\ &= \begin{pmatrix} c_{12} c_{13} & s_{12} c_{13} & s_{13} e^{-i\delta_{13}} \\ -s_{12} c_{23} - c_{12} s_{23} s_{13} e^{i\delta_{13}} & c_{12} c_{23} - s_{12} s_{23} s_{13} e^{i\delta_{13}} & s_{23} c_{13} \\ s_{12} s_{23} - c_{12} c_{23} s_{13} e^{i\delta_{13}} & -c_{12} s_{23} - s_{12} c_{23} s_{13} e^{i\delta_{13}} & c_{23} c_{13} \end{pmatrix}. \end{aligned} \quad (2.23)$$

Since the phase δ is the only complex term, it encodes the entirety of the CP violation in the quark transitions.

As this parameterisation is rather unwieldy, another common parameterisation, the Wolfenstein parameterisation, gives these elements as

$$V_{\text{CKM}} = \begin{pmatrix} 1 - \frac{1}{2}\lambda^2 & \lambda & A\lambda^3(\rho - i\eta) \\ -\lambda & 1 - \frac{1}{2}\lambda^2 & A\lambda^2 \\ A\lambda^3(1 - \rho - i\eta) & -A\lambda^2 & 1 \end{pmatrix} + \mathcal{O}(\lambda^4), \quad (2.24)$$

2 Heavy flavour physics in the Standard Model

where $\lambda = s_{12}$, $A = s_{23}/s_{12}^2$, $\rho + i\eta = s_{13}e^{-i\delta_{13}}/s_{12}s_{23}$ (with $\rho, \eta \in \mathbb{R}$), are the Wolfenstein parameters [20]. This is an expansion in powers of λ , which shows that the diagonal elements are closest to 1 (Cabibbo-favoured) and that the off-diagonal (Cabibbo-suppressed) elements are then significantly smaller, with V_{ub} and V_{td} as the smallest. The latest experimental averages of these parameters are $\lambda = 0.22501 \pm 0.00068$, $A = 0.826_{-0.015}^{+0.016}$, $\rho = 0.1591 \pm 0.0094$, and $\eta = 0.3523_{-0.0071}^{+0.0073}$ [5].

The CKM matrix is, by definition, unitary, *i.e.*, $V_{\text{CKM}}V_{\text{CKM}}^\dagger = I_3$. This can be written as two constraints in terms of the elements:

$$\sum_j V_{ij}V_{ij}^* = \sum_j V_{ij}V_{ij}^* = 1, \quad (2.25a)$$

$$\sum_k V_{ik}V_{jk}^* = 0 \quad \text{for } i \neq j, \quad (2.25b)$$

constraining the diagonal and off-diagonal elements, respectively. The latter is particularly interesting: the sum of three complex numbers to 0 can be considered a closed triangle in the complex space. Therefore, Eq. 2.25b represents six triangles, one for each i, j combination. The area of these triangles is given by half of the Jarlskog invariant, J , written as

$$J = \left| \text{Im} \left\{ V_{ij}V_{jk}V_{il}^*V_{kj}^* \right\} \right| = c_{12}c_{13}^2c_{23}s_{12}s_{13}s_{23} \sin \delta. \quad (2.26)$$

Since this is proportional to $\sin \delta$, it is considered a measure of the scale of CP violation in the SM, estimated to be $J \sim 3 \times 10^{-5}$ [21]. This also implies that three generations of quarks are required for CP violation to be present.

Considering the case of $i = d$ and $j = b$, the triangle can be scaled by $|V_{cd}V_{cb}^*|$ such that its points lie on $(0, 0)$, $(0, 1)$, and $(\bar{\rho}, \bar{\eta})$, wherein $\bar{\rho}$ and $\bar{\eta}$ are the generalised Wolfenstein parameters $\bar{\rho} = \rho \sum_n (-1)^n \lambda^n / n!$ and $\bar{\eta} = \eta \sum_n (-1)^n \lambda^n / n!$, such that the angles of the triangle are given by

$$\alpha = \arg \left(-\frac{V_{td}V_{tb}^*}{V_{ud}V_{ub}^*} \right), \quad (2.27a)$$

$$\beta = \arg \left(-\frac{V_{cd}V_{cb}^*}{V_{ud}V_{ub}^*} \right), \quad (2.27b)$$

$$\gamma = \arg \left(-\frac{V_{cd}V_{cb}^*}{V_{td}V_{tb}^*} \right). \quad (2.27c)$$

This triangle, and accompanying experimental results (as of September 2023), are shown in Fig. 2.2. Measurements of the parameters of the CKM unitarity triangle provide a null test of the SM in the closure of the triangle, and provide access to the scale of CP violation [22].

In contrast to charged currents, neutral currents do not involve a change of electric charge and cannot facilitate a change in flavour at tree-level. At higher orders, flavour-changing neutral currents (FCNCs) are possible in the SM via the Glashow-Iliopoulos-Maiani (GIM) mechanism, wherein W^\pm bosons are exchanged in one-loop box-diagrams containing a virtual quark [24]. As the possible box-diagrams contain elements of the CKM matrix on either side of the diagonal, which are of similar magnitude but opposite sign, the resulting interference between these loop diagrams is destructive and leads to a near-total cancellation. This cancellation, known as GIM suppression, is proportional to the difference in m_q^2/m_W^2 for the involved virtual quarks q .




Figure 2.2: CKM unitarity triangle defined by the angles in Eqs. 2.27a-2.27c, with the constraints imposed by relevant measurements overlaid, taken from Ref. [23].

The couplings of the W^\pm , Z and γ to charged leptons are independent of the lepton flavour; processes involving each flavour of lepton should, accounting for differences in their masses, proceed at the same rate. This property is known as lepton flavour universality and has been studied extensively, as discussed in Section 2.4.

2.2.4 Phenomenology of heavy flavour physics

The properties and decays of hadrons containing the heavy b - and c -quarks¹ provide many ways to stress test the SM and search for new physics. Before embarking on searches for physics beyond the SM (BSM), as is discussed in Section 2.3, we must first understand how these hadrons behave.

This work focuses primarily on hadrons containing a b -quark. As the b is always the heaviest quark in a b -hadron, the dominant decays of a b -hadron are those which remove the b -quark. Since, at tree-level, only the weak charged-current interaction can change quark flavours, b -hadrons decay predominantly via the weak interaction and hence have a longer lifetime than hadrons which decay primarily via the strong interaction. This lifetime is typically around $\mathcal{O}(1 \text{ ps})$, which, at energies of $\sim 10 \text{ TeV}$, corresponds to a flight distance (FD) of $\sim 1 \text{ cm}$. The detachment from the point of production, the primary vertex (PV), also causes the decay products to appear to originate from away from the PV. This is quantified by the impact parameter (IP), the closest distance between the PV and the trajectory of a given particle. The topology of an archetypal decay of a b -hadron to a c -hadron is shown in Fig. 2.3, with the FDs of both hadrons and the IP of the c -hadron annotated.

The quantum states of neutral mesons containing b -quarks exhibit oscillation between their flavour eigenstates, *i.e.*, between $B_{(s)}^0$ and $\bar{B}_{(s)}^0$, over periods of a similar scale to their respective lifetimes [27, 28]. This is in contrast to neutral kaons or c -mesons, in which the smaller difference in mass of the constituent quarks results in a heavier suppression by the GIM mechanism [24], resulting in much longer oscillation periods.

For a neutral meson M containing $q\bar{q}'$, the flavour eigenstates of the meson and its antimeson counterpart are $|M\rangle$ and $|\bar{M}\rangle$, respectively. The mass eigenstates of the

¹The t -quark is excluded in most definitions of heavy flavour physics, as it is too heavy, and hence short-lived, to form bound states. Recent measurements by the ATLAS and CMS experiments suggest the existence of a quasi-bound $t\bar{t}$ state, though this is produced extremely rarely [25, 26].

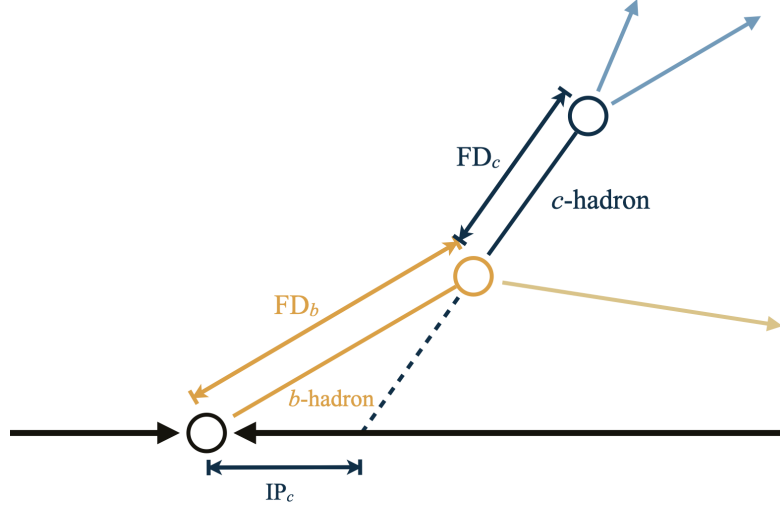


Figure 2.3: Schematic of the topology of a b -hadron decaying to a c -hadron. Distances not drawn to scale.

meson, from which its time-evolution is defined, can be written as a linear combination of these,

$$|M_{1,2}\rangle = p|M\rangle - q|\bar{M}\rangle, \quad (2.28)$$

with corresponding masses $m_{1,2}$ and decay widths $\Gamma_{1,2}$, where $|p|^2 + |q|^2 = 1$. From the Schrödinger equation, these states derive a time-evolution of

$$|M_{1,2}(t)\rangle = e^{-(im_{1,2} + \Gamma_{1,2}/2)t} |M_{1,2}(0)\rangle. \quad (2.29)$$

Incorporating this into the inverse of Eq. 2.28, the time-evolution of the flavour eigenstates can be written as

$$|M(t)\rangle = \frac{1}{2p} \left(\left(e^{-(im_1 + \Gamma_1/2)t} + e^{-(im_2 + \Gamma_2/2)t} \right) p|M(0)\rangle + \left(e^{-(im_1 + \Gamma_1/2)t} - e^{-(im_2 + \Gamma_2/2)t} \right) q|\bar{M}(0)\rangle \right). \quad (2.30)$$

The time evolution can also be described in terms of the means and differences in masses ($m = (m_1 + m_2)/2$ and $\Delta m = m_2 - m_1$) and decay widths ($\Gamma = (\Gamma_1 + \Gamma_2)/2$ and $\Delta\Gamma = \Gamma_2 - \Gamma_1$). These can be reparameterised further, introducing the terms $x = \Delta m/\Gamma$ and $y = \Delta\Gamma/2\Gamma$, such that the time-dependent prefactors of Eq. 2.30 can be written as

$$f_{\pm}(t) = \frac{1}{2} e^{-(im + \Gamma/2)t} \left(e^{(ix + y)\Gamma t/2} \pm e^{-(ix + y)\Gamma t/2} \right). \quad (2.31)$$

This is a handy condensation, as the square magnitudes of $f_{\pm}(t)$ are then simply

$$|f_{\pm}|^2 = \frac{1}{2} e^{-\Gamma t} (\cosh(y\Gamma t) \pm \cos(x\Gamma t)). \quad (2.32)$$

The probabilities that a meson starts in flavour eigenstate $|M\rangle$ and either remains in this state or has transitioned to $|\bar{M}\rangle$ at a time, t , are given by

$$P(M \rightarrow M, t) = |\langle M|M(t)\rangle|^2 = |f_+(t)|^2, \quad (2.33a)$$

$$P(M \rightarrow \bar{M}, t) = \left| \frac{q}{p} \right|^2 |f_-(t)|^2. \quad (2.33b)$$

Similarly for $|\bar{M}\rangle$, the probabilities are

$$P(\bar{M} \rightarrow \bar{M}, t) = P(M \rightarrow M, t), \quad (2.34a)$$

$$P(\bar{M} \rightarrow M, t) = \left| \frac{p}{q} \right|^2 |f_-(t)|^2. \quad (2.34b)$$

The flavour of the meson oscillates with a regular period, defined primarily by the relative difference in the masses (x) of the two eigenstates; for a smaller mass difference, the GIM suppression of the mixing amplitudes is greater and the period longer. For charm mesons, such as D^0 , where x is very small (measured to be $x \sim 4 \times 10^{-3}$ [5]), oscillations occur very slowly, with effectively all mesons unable to complete a full period before decaying. In contrast, for B_s^0 mesons, x is very large (measured to be $x = 27.01 \pm 0.10$ [5]), with many mesons completing multiple oscillation periods before decaying. The measurement of these oscillations in $B_s^0 \rightarrow D_s^- \pi^+$ decays by LHCb is shown in Fig. 2.4.

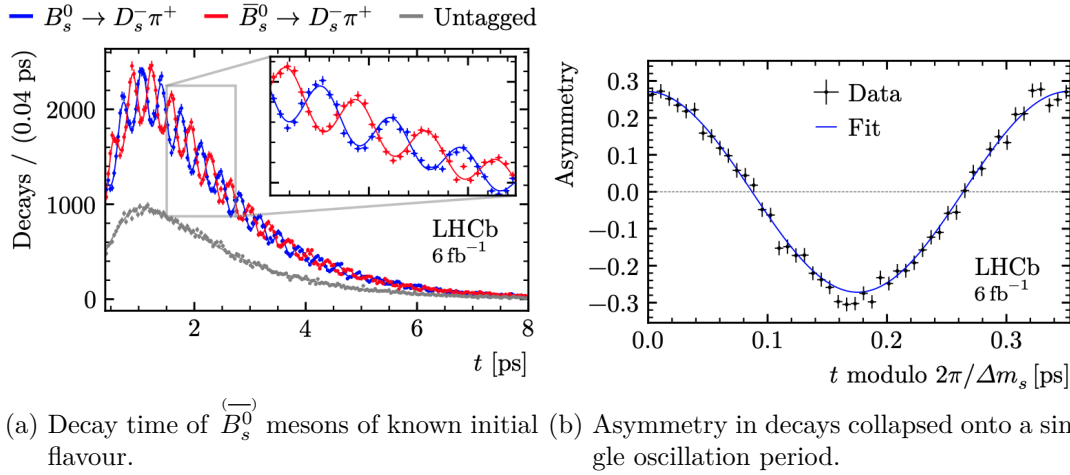


Figure 2.4: Mixing of B_s^0 and \bar{B}_s^0 mesons as measured by LHCb in $B_s^0 \rightarrow D_s^- \pi^+$ decays in Run 2, taken from Ref. [29].

In addition to the mass eigenstates, CP eigenstates can be formed from the flavour eigenstates:

$$|M_{\pm}(t)\rangle = \frac{1}{\sqrt{2}} (|M(t)\rangle \pm |\bar{M}(t)\rangle), \quad (2.35)$$

such that $CP|M_{\pm}\rangle = \pm|M_{\pm}\rangle$. These provide the rates at which $|\bar{M}\rangle$ decay to either final state f or \bar{f} . The CP asymmetry can then be defined as the ratio of difference to sum between the rates of decays to f and \bar{f} :

$$A_{CP}(M \rightarrow f) = \frac{[\Gamma(M(t) \rightarrow \bar{f}) + \Gamma(\bar{M}(t) \rightarrow \bar{f})] - [\Gamma(M(t) \rightarrow f) + \Gamma(\bar{M}(t) \rightarrow f)]}{[\Gamma(M(t) \rightarrow \bar{f}) + \Gamma(\bar{M}(t) \rightarrow \bar{f})] + [\Gamma(M(t) \rightarrow f) + \Gamma(\bar{M}(t) \rightarrow f)]}. \quad (2.36)$$

CP violation in $B_{(s)}^0$ decays can arise from three different sources. Firstly, direct CP violation arises when the interference between different decay amplitudes causes $\Gamma(M \rightarrow f) \neq \Gamma(\bar{M} \rightarrow \bar{f})$, *e.g.*, from the tree-level and penguin diagrams of $B^0 \rightarrow K^+ \pi^-$ decays

as studied in Ref. [30]. Secondly, indirect CP violation arises through the mixing of neutral mesons, wherein CP violation is observed if $\Gamma(M \rightarrow \bar{M}) \neq \Gamma(\bar{M} \rightarrow M)$, as studied in $B^0 \rightarrow D^{*-}\mu^+\nu_\mu$ decays in Ref. [31]. Finally, interference CP violation can arise when M and \bar{M} both decay to f , as M may oscillate to \bar{M} and then decay to f . This additional amplitude then interferes with the direct decay $M \rightarrow f$, such that $\Gamma(M \rightarrow f) \neq \Gamma(M \rightarrow \bar{M} \rightarrow f)$, *e.g.*, in the decay $B^0 \rightarrow J/\psi K_S^0$ studied in Ref. [32].

Considering first direct CP violation, we take a flavour-specific process $M \rightarrow f$, *i.e.*, a process in which M decays to f but not \bar{f} and \bar{M} decays to \bar{f} but not f , with at least two contributing amplitudes, say $\mathcal{A}_1(M \rightarrow f) = A_1 e^{i\delta_1} e^{i\phi_1}$ and $\mathcal{A}_2(M \rightarrow f) = A_2 e^{i\delta_2} e^{i\phi_2}$, wherein A_i , δ_i , ϕ_i are the magnitude, strong phase and weak phase of amplitude \mathcal{A}_i [33]. Their sum can be written in terms of the relative magnitude and phases, $r = A_2/A_1$, $\delta = \delta_2 - \delta_1$ and $\phi = \phi_2 - \phi_1$, as

$$\mathcal{A}(M \rightarrow f) = A \left(1 + r e^{i\delta} e^{i\phi} \right). \quad (2.37)$$

Assuming that the equivalent amplitude for $\bar{M} \rightarrow \bar{f}$ differs only by the sign of ϕ , then the CP asymmetry is given, per Eq. 2.36 by

$$\begin{aligned} A_{CP}(M \rightarrow f) &= \frac{\mathcal{A}_1^2 |1 + r e^{i\delta} e^{-i\phi}|^2 - \mathcal{A}_1^2 |1 + r e^{i\delta} e^{i\phi}|^2}{\mathcal{A}_1^2 |1 + r e^{i\delta} e^{-i\phi}|^2 + \mathcal{A}_1^2 |1 + r e^{i\delta} e^{i\phi}|^2} \\ &\approx 2r \sin(\phi) \sin(\delta), \end{aligned} \quad (2.38)$$

for small values of r , δ and ϕ . Provided that the two amplitudes differ in both their strong and weak phases, direct CP violation will be induced.

Indirect CP violation is rather different, instead arising from differences in the processes $M \rightarrow \bar{M}$ and $\bar{M} \rightarrow M$ [27]. We consider again a flavour-specific process: if the initial state is observed to be M (*e.g.*, via techniques such as flavour tagging described in Section 3.2.3) and the final state observed as \bar{f} , then we know that the decay must have proceeded as $M \rightarrow \bar{M} \rightarrow \bar{f}$. Making use of the probabilities in Eqs. 2.33b and 2.34b, the decay rates involved in the asymmetry are:

$$\Gamma(M \rightarrow \bar{f}, t) = \Gamma(\bar{M} \rightarrow M \rightarrow \bar{f}, t) \propto \left| \frac{p}{q} \right|^2 |f_-(t)|^2, \quad (2.39a)$$

$$\Gamma(\bar{M} \rightarrow f, t) = \Gamma(\bar{M} \rightarrow M \rightarrow f, t) \propto \left| \frac{q}{p} \right|^2 |f_-(t)|^2. \quad (2.39b)$$

The asymmetry between these decay rates is then given by

$$A_{CP} = \frac{|q/p|^2 - |p/q|^2}{|q/p|^2 + |p/q|^2}, \quad (2.40)$$

which is notably time-independent, as the terms in $f_-(t)$ cancel. If $|q/p| \neq 1$, CP violation is present in mixing.

Finally, considering interference CP violation, for a process in which both M and \bar{M} can decay to a CP invariant final state f , the interference of decays proceeding directly ($M \rightarrow f$) and via mixing ($M \rightarrow \bar{M} \rightarrow f$) result in CP violation. Adapting Eq. 2.30, the amplitudes can be written as

$$\mathcal{A}(M(t) \rightarrow f) = \mathcal{A}_f f_+(t) + \bar{\mathcal{A}}_f \frac{q}{p} f_-(t) \quad (2.41a)$$

$$\mathcal{A}(\overline{M}(t) \rightarrow f) = \overline{\mathcal{A}}_f f_+(t) + \mathcal{A}_f \frac{p}{q} f_-(t). \quad (2.41b)$$

Applying a similar treatment to that between Eqs. 2.31 and 2.32, and introducing the factor $\lambda_f = \frac{q}{p} \frac{\overline{\mathcal{A}}_f}{\mathcal{A}_f}$ the asymmetry induced by the interference of these amplitudes is given by

$$\begin{aligned} A_{CP}(t) &= \frac{\Gamma(\overline{M}(t) \rightarrow f) - \Gamma(M(t) \rightarrow f)}{\Gamma(\overline{M}(t) \rightarrow f) + \Gamma(M(t) \rightarrow f)} \\ &= \frac{2\text{Im}\{\lambda_f\} \sin(x\Gamma t) - (1 - |\lambda_f|^2) \cos(x\Gamma t)}{2\text{Re}\{\lambda_f\} \sinh(y\Gamma t) + (1 + |\lambda_f|^2) \cosh(y\Gamma t)}. \end{aligned} \quad (2.42)$$

For B^0 mesons, where the difference in the decay widths, and, by extension, y , is effectively 0, this reduces further to

$$A_{CP}(t) = S_f \sin(x\Gamma t) - C_f \cos(x\Gamma t), \quad (2.43)$$

where the scale of interference CP violation is parameterised by [34]

$$S_f = \frac{2\text{Im}\{\lambda_f\}}{1 + |\lambda_f|^2}, \quad (2.44a)$$

$$C_f = \frac{1 - |\lambda_f|^2}{1 + |\lambda_f|^2} \quad (2.44b)$$

The variety of observables with which CP violation in the b sector can be studied provides numerous routes to probe the SM. Measurements of these observables, such as in Chapter 4, rely on large samples of b -hadron decays. Fortunately, many decays in which these observables can be studied proceed by charged currents, *i.e.*, they contain $b \rightarrow c$ and $b \rightarrow u$ transitions, and hence proceed with branching fractions upwards of $\mathcal{O}(10^{-4})$.

A final set of b -hadron processes to consider are decays containing a $b \rightarrow s$ or $b \rightarrow d$ transition. These transitions involve FCNCs, which, as discussed in Section 2.2.3, are forbidden at tree-level in the SM. Instead, they must proceed via a loop, as the b -quark must first be transition to a virtual up-type quark before then transitioning to a s or d , for example, *e.g.*, as shown in Fig. 2.5 for the transition $b \rightarrow s \ell^+ \ell^-$.




Figure 2.5: Leading order Feynman diagrams for $b \rightarrow s \ell^+ \ell^-$ transitions in the SM.

The virtual quark is most likely a t as, V_{tb} is several orders of magnitude larger than V_{ub} and V_{cb} , as demonstrated in Eq. 2.24, where V_{ub} and V_{cb} are suppressed by λ^3 and

λ^2 , respectively. As this t -quark must then transition to the s or d , the CKM element V_{sb}^* or V_{db}^* is introduced, further suppresses the amplitude. Additional suppression of such processes is also induced by the GIM mechanism. The CKM and GIM suppression result in extremely small amplitudes, such that processes containing these transitions are referred to as rare decays. For example, the decays $B^+ \rightarrow K^+ \ell^+ \ell^-$, $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ and $B_s^0 \rightarrow \phi \ell^+ \ell^-$ have branching fractions measured to be $(4.53 \pm 0.35) \times 10^{-7}$, $(9.4 \pm 0.5) \times 10^{-7}$ and $(8.3 \pm 0.4) \times 10^{-7}$, respectively [5].

For the SM to be probed by studying the nature of b -hadron decays described above, precise theoretical predictions must be made of the expected values of these properties. Many frameworks exist for computing such predictions; however, those which have proven most successful use the unique properties of b -hadrons, such as their relatively large mass, to their advantage.

2.2.5 Effective field theories for heavy flavour phenomenology

Effective field theories (EFTs) provide approximate descriptions of a field theory, constructing an effective Hamiltonian containing only the aspects of the theory which are relevant at a given energy scale. As the heavy quark in a b -hadron is much heavier than the other (anti)quark(s), an EFT describing the behaviour of b -hadrons, known as the Heavy Quark Effective Theory (HQET), can be defined from the scale of the b -quark mass, such that quantities are expressed as series expansions in powers of $1/m_b$. These series can then be used in perturbative calculations to obtain predictions of heavy flavour observables, such as hadronic masses and decay rates. The applications of HQET to the two families of b -hadrons decays on which this thesis focuses are considered: non-leptonic decays containing a $b \rightarrow \bar{c}uq$ transition (specifically $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays) and rare decays containing a $b \rightarrow s \ell^+ \ell^-$ transition.

Predictions of observables for $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays, require careful treatment of contributions from QCD, as both the initial and final states are fully hadronic. Perturbative calculations require quantities to be expressed as series of diminishing small terms. This is not the case for the long-distance contributions from QCD, *i.e.*, the behaviour of strongly coupled constituent quarks and gluons, hence these contributions cannot be neatly considered together with the short-distance contributions from the weak interactions in which we are interested. The framework to disentangle these contributions is known as QCD factorisation (QCDF) [35], wherein the fact that $\Lambda_{\text{QCD}}/m_b \ll 1$ is used to form an expansion in terms of Λ_{QCD}/m_b which cleanly separates the scales of the QCD contributions [36].

In particular, two $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays, $B^0 \rightarrow D^- K^+$ and $B_s^0 \rightarrow D_s^- \pi^+$, are considered golden modes for QCDF, as these flavour-specific decays proceed (at leading orders) solely at tree-level, avoiding non-factorisable contributions from penguin and topologies [37]. QCDF has thus been used extensively to predict observables such as the branching fractions discussed in Section 2.4 and the ratios of b -hadron fragmentation fractions [38].

In the case of decays containing $b \rightarrow s \ell^+ \ell^-$ transitions, the leptons in the final state cannot interact with the b and s via the strong interaction; aside from in the internal meson dynamics, only the weak interaction is at play. Neglecting small corrections from the heavily CKM suppressed contributions in which the virtual quark is a u or c , the

effective Hamiltonian describing $b \rightarrow s\ell^+\ell^-$ transitions can be written as an operator-product expansion:

$$\mathcal{H}_{\text{eff}} = -\frac{4G_F}{\sqrt{2}} V_{tb} V_{ts}^* \sum_i (\mathcal{C}_i \mathcal{O}_i + \mathcal{C}'_i \mathcal{O}'_i), \quad (2.45)$$

where \mathcal{O}_i are operators describing the possible left-handed long-distance interactions of the four fermions, \mathcal{C}_i are their corresponding Wilson coefficients (WCs), and \mathcal{O}'_i and \mathcal{C}'_i are their right-handed equivalents [39, 40]. These operators correspond directly to different Feynman diagrams contributing to $b \rightarrow s\ell^+\ell^-$ amplitudes, *e.g.*, the EM penguin, EW penguin, and box-loop, which corresponding to operators $\mathcal{O}_7^{(\prime)}$, $\mathcal{O}_9^{(\prime)}$ and $\mathcal{O}_{10}^{(\prime)}$. At the energy scale of the b -quark mass, these are given by

$$\mathcal{O}_7^{(\prime)} = \frac{e}{16\pi^2} m_b \left(\bar{s} \sigma_{\mu\nu} P_{R(L)} b \right) F^{\mu\nu}, \quad (2.46a)$$

$$\mathcal{O}_9^{(\prime)} = \frac{e^2}{16\pi^2} \left(\bar{s} \gamma_\mu P_{L(R)} b \right) \left(\bar{\ell} \gamma^\mu \ell \right), \quad (2.46b)$$

$$\mathcal{O}_{10}^{(\prime)} = \frac{e^2}{16\pi^2} \left(\bar{s} \gamma_\mu P_{L(R)} b \right) \left(\bar{\ell} \gamma^\mu \gamma_5 \ell \right), \quad (2.46c)$$

in which $P_L = (1 - \gamma_5)/2$ and $P_R = (1 + \gamma_5)/2$ [39].

Many observables provide access to the three WCs listed above: branching fractions of $b \rightarrow s\ell^+\ell^-$ processes relate $\mathcal{C}_{7,9,10}$ [41], angular observables in $B^0 \rightarrow K^{*0}\ell^+\ell^-$ decays relate to $\mathcal{C}_7^{(\prime)}$ [41], and the ratios of branching ratios between dielectron and dimuon modes probe $\mathcal{C}_{9,10}^{(\prime)}$ [42]. The WCs can be measured and compared to their theoretical predictions: contributions from NP could reduce or enhance one or more WC. The WCs thus give an indication of both the scale and topology of such contributions.

2.3 Looking beyond the SM

The SM has proven thus far to be an extremely robust, predictive model. The existence of the gluon, W^\pm , Z , Higgs boson, τ and the c -, b - and t -quarks, were all predicted by the SM prior to their discoveries. Similarly, the SM has survived countless tests, including direct searches and null tests of its predicted phenomena. However, when we step back and look at fundamental physics as a whole, we see many phenomena which cannot be explained by the SM, a handful of which are discussed below.

Firstly, are the phenomena and properties of the universe on the largest scales, supported by extensive cosmological evidence such as the angular power spectrum of the cosmic microwave background [43]. This evidence affirms general relativity, indicates that the expansion of the universe is accelerating, and supports the Lambda cold dark matter model in which the universe contains dark energy and dark matter in addition to ordinary matter. The SM provides no description of gravity or a viable cold dark matter candidate. Additionally, significantly more (baryonic) matter than antimatter is observed in the universe than can be accounted for based on the CP violation of the SM [44, 45].

The description of neutrinos in the SM has also been proven to be incomplete. The oscillation of neutrinos between flavours, first directly observed by the Sudbury National Observatory in 2001 [46] and further confirmed by later experiments, is only possible if at least two of the neutrinos are massive. In the SM, all three neutrinos are massless, thus an extension to the SM must be required, provided by the Pontecorvo-Maki-Nakagawa-Sakata matrix, which is analogous to the CKM matrix in that each neutrino obtains

2 Heavy flavour physics in the Standard Model

a mass eigenstate, $m_{1,2,3}$. Neither the form of the neutrino masses (whether these are given by Dirac- or Majorana-like terms [47]) or their ordering [48] are specified by the SM.

In addition to these discrepancies, the electroweak scale ($\mathcal{O}(100 \text{ GeV})$) is much smaller than the Planck scale ($\mathcal{O}(10^{19} \text{ GeV})$) without an apparent explanation, known as the hierarchy problem. In particular, the corrections to the bare Higgs mass arising from vacuum polarisation are quadratically dependent on the energy scale [49]. If the SM is valid at the Planck scale, then for the Higgs to obtain its observed mass, either the bare Higgs mass must be extremely finely-tuned, *i.e.*, to within 10^{-34} , or additional contributions must otherwise provide a cancellation of the quantum corrections [50].

It is clear that the SM is incomplete and that some form of physics lies beyond the SM. We thus turn our attention to the ways in which NP, particularly new particles, can be sought. Searches for NP fall under two categories: direct searches which seek new particles through their production, and indirect searches which study observables affected by the existence of new particles. Direct searches provide the clearest possible evidence for the existence of new particles; however, they require a system with sufficient energy to produce the new particles, *i.e.*, the maximum collision energy of collider experiments. In contrast, indirect searches are inherently less sensitive to NP than direct searches; however, as the new particles need only exist as virtual particles, indirect searches can probe much higher energy scales.

As discussed in Section 2.2.4, the decays of b -hadrons provide a range of observables with which to test the SM; the frameworks discussed in Section 2.2.5 provide precise SM predictions against which experimental observations can be compared. The simplest of such observables to study are the branching fractions of decays: BSM particles may introduce additional amplitudes to a given decay, which interfere with the SM decay amplitude to result in a branching fraction larger or smaller than predicted by the SM. If these additional amplitudes have weak and strong phases differing from the SM amplitude then such BSM particles would also induce direct CP violation per Eq. 2.38, which can be measured in decays containing $b \rightarrow c\bar{u}q$ transitions as these proceed relatively abundantly. In the flavour-specific QCDF golden modes of $B^0 \rightarrow D^- K^+$ and $B_s^0 \rightarrow D_s^- \pi^+$, where the decays are described by a single amplitude, the addition of such an additional amplitude would modify the asymmetry A_{fs}^q as

$$A_{\text{fs}}^q \approx a_{\text{fs}}^q - A_{\text{dir}}^q, \quad (2.47)$$

where a_{fs}^q is the CP asymmetry from mixing, which should be equal to a_{sl}^q , the CP asymmetry from mixing measured in semi-leptonic decays of the $D_{(s)}^-$. If A_{fs}^q were to be found to be inconsistent with a_{sl}^q , this would be a clear indication of the presence of NP. Such a test is described in Chapter 4. Similar modifications of CP violation in other channels may also induce deviations in tests of CKM unitarity.

The branching fractions of rare decays, *e.g.*, those proceeding by FCNCs such as $b \rightarrow s\ell^+\ell^-$, are those most sensitive to NP. Since the branching fractions of these are so small, *e.g.*, $\mathcal{B}(B^0 \rightarrow K^* \mu^+ \mu^-) = (9.4 \pm 0.5) \times 10^{-8}$ [5], even a small additional amplitude could induce a significant excess. The dependence of branching fractions on the angular properties of the $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ system can also be modified by NP contributions. Finally, the ratios of branching fractions between decays containing different leptons, *e.g.*, the ratio $\mathcal{B}(B^+ \rightarrow K^+ \mu^+ \mu^-) / \mathcal{B}(B^+ \rightarrow K^+ e^+ e^-)$ known as R_K , are modified by some BSM models in which LFU is violated, *i.e.*, in which new particles couple to each lepton flavour with a different strength. LFU ratios are particularly interesting observables as

the SM predictions are very clean and the form factors, on which the uncertainties are large, cancel out in the ratio. Studies validating the Run 3 LHCb detector (described in Chapter 5) ahead of measurements of such ratios using Run 3 data are detailed in Chapter 6. As discussed in Section 2.2.5, all of these quantities provide access to the WCs $C_{7,9,10}^{(\prime)}$, allowing for comparisons of combinations of observables against SM predictions.

2.4 Current experimental status

Having defined observables in which NP can be sought in the decays of b -hadrons, the status of measuring these observables is considered. In particular, this section focuses on the $b \rightarrow c\bar{u}q$ anomaly and anomalies in tests of LFU, which motivate the analysis described in Chapter 4 and measurement described in Chapter 6, respectively.

2.4.1 The $b \rightarrow c\bar{u}q$ anomaly

In recent years, measurements and theoretical predictions of the branching fractions of $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays have exhibited growing tensions [38]. These tensions, summarised in Fig. 2.6, all involve deficits with respect to the predicted values, which are all of similar relative scales across the family of decays. Since these decays all contain a $b \rightarrow c\bar{u}q$ transition, this is known as the $b \rightarrow c\bar{u}q$ anomaly.

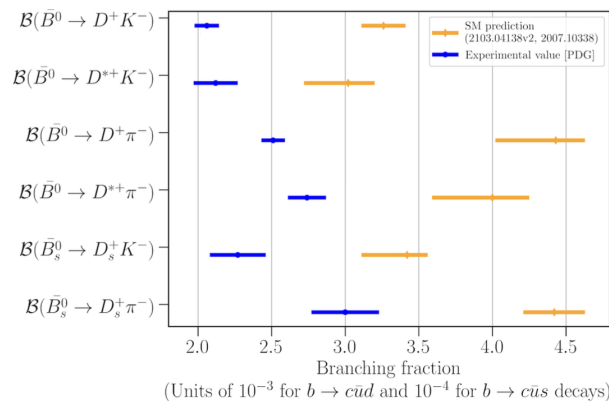


Figure 2.6: Overview of experimental (blue) and theoretical (yellow) averages of branching fractions of $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays, courtesy of Nicole Skidmore.

As described in Section 2.2.5, QCDF (with which these branching fractions are predicted) is widely considered a reliable and robust framework; such tensions come as a surprise and we should consider an alternative explanation for these. As in Section 2.3, an additional contribution to the decay amplitudes from a generic NP model could result in such a tension. However, branching fractions can be difficult to measure precisely as they require a deep understanding of accelerator, detector and selection effects, and are thus typically cautiously compared to theoretical predictions. If NP is responsible for these tensions, then other observables may be sensitive to this [51]. We therefore turn to the CP asymmetry, as per Eq. 2.47.

In Chapter 4, a direct contribution to the CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays, A_{FB}^s , is studied in LHCb pp data recorded in Run 2 of the LHC.

2.4.2 Lepton flavour anomalies

As described in Section 2.3, NP contributions which couple more strongly to different flavours of leptons may induce excesses or deficits in the branching fractions of decays containing $b \rightarrow s\ell^+\ell^-$ transitions, and values of LFU ratios inconsistent with unity. In addition to providing clean theoretical predictions, per Section 2.2.5, both the displaced primary vertex from the b -hadron decay and the dilepton ($\ell^+\ell^-$) final state are relatively clean signatures, making such decays experimentally advantageous to study. Dedicated selection algorithms can be derived which target these signatures, such as those discussed in Section 5.2.1.

Branching fractions and LFU ratios spanning a broad range of dilepton signatures have been measured and compared extensively against SM predictions, as shown in Fig. 2.7.

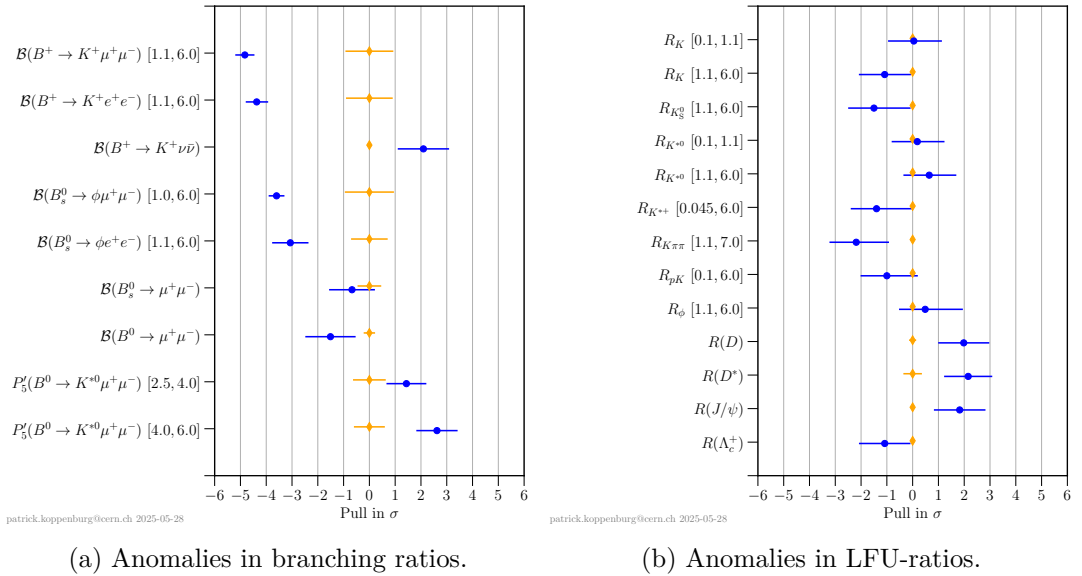


Figure 2.7: Summaries of the current state of $b \rightarrow s\ell^+\ell^-$ anomalies taken from Ref. [52].

The branching fractions exhibit tensions with SM predictions across the many channels studied, in places up to almost 5σ . As for the $b \rightarrow c\bar{u}q$ anomaly, tensions in these branching fractions are treated with caution; however, the scale of the tension leaves plenty of room for NP contributions. Since the SM predictions of LFU ratios are much cleaner, a large tension between prediction and measurement in any one of these would be conclusive evidence of NP. Small tensions in the LFU ratios are present across those measured, though only up to $\sim 2\sigma$. The largest of these are seen in $R(D)$ and $R(D^*)$, which when compared together against the equivalent SM predictions, demonstrate a greater tension of 3.14σ [53].

In Chapter 4, the LFU ratio of $B^+ \rightarrow J/\psi(\ell^+\ell^-)K^+$ decays, $r_{J/\psi}$, a prerequisite to R_K , is measured in LHCb pp data recorded in 2024.

3 The LHCb experiment during Run 2 of the LHC

The Large Hadron Collider beauty (LHCb) experiment is one of four major experiments at the Large Hadron Collider (LHC), designed for studying heavy flavour physics. This chapter introduces the LHC and the LHCb detector as operated between 2015 and 2018. The high-quality tracking and identification of charged particles by the LHCb detector are also introduced. These make LHCb a uniquely well-equipped observatory for measurements of heavy flavour processes, such as those described in Chapters 4 & 6. The trigger systems employed to process the immense data recorded by the LHCb detector are also discussed, ahead of the detailed discussion of their upgrade in Chapter 5.

3.1 The Large Hadron Collider

The LHC, shown in Fig. 3.1, was designed to collide beams of protons and heavy ions such as lead, though this thesis will focus solely on the pp configuration. Large experiments are placed at each of four interaction points (IPs), where the beams are brought to collide: A Toroidal LHC Apparatus (ATLAS), A Large Ion Collision Experiment (ALICE), Compact Muon Solenoid (CMS), and LHCb. The LHC has operated for three data-taking runs: Run 1 (2009-2013), Run 2 (2015-2018), and Run 3 (2022-present, expected mid-2026). Additional runs are also planned to take place into the 2040s.

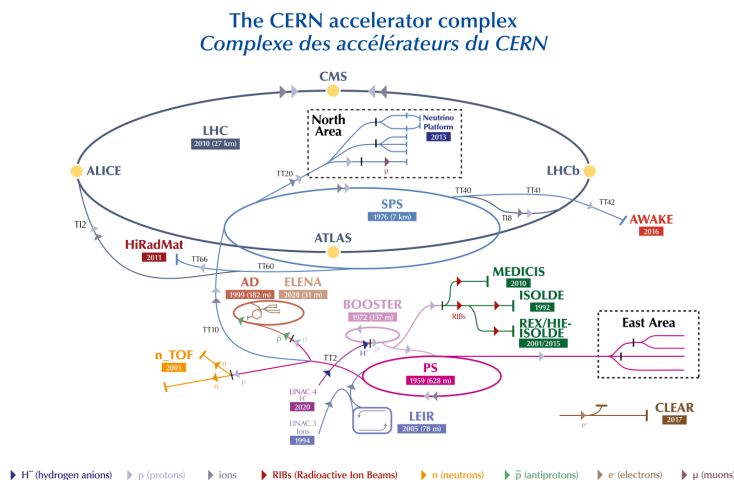


Figure 3.1: Schematic diagram of the LHC and accompanying accelerator facilities, sourced from Ref. [54].

Located on the French-Swiss border near Geneva, the LHC occupies a 27 km tunnel, designed to accelerate and collide counter-rotating proton beams at centre-of-mass energies of up to 14 TeV at an instantaneous luminosity of up to $2 \times 10^{34} \text{ cm}^{-2} \text{ s}^{-1}$. This luminosity is achieved by operating with beams of 2808 proton bunches, each containing 1.1×10^{11} protons, which are injected in *fills*. The bunches are separated such that collisions occur every 25 ns, *i.e.*, at 40 MHz. The protons, which are provided by the ionisation of hydrogen, are accelerated to their nominal collision energy by a chain of accelerators, each operating in a different energy regime.

3 The LHCb experiment during Run 2 of the LHC

Protons are first accelerated to 160 MeV by Linear Accelerator 4 [55], then by the circular accelerators of the BOOSTER to 1.4 GeV, the Proton Synchrotron to 25 GeV, and the Super Proton Synchrotron to 450 GeV [56]. These protons are then injected into the LHC, which has 8 bent and 8 straight segments: bent segments employ dipole magnets (1232 in total) to bend the beams; straight segments have a variety of purposes. Of the 8 straight segments, one employs radio-frequency cavities to accelerate each beams to 6.5 TeV, another is responsible for dumping the beams, two segments employ quadrupole magnets to collimate the beams, and the final four contain the IPs.

The IPs are the locations of the four LHC experiments. At IP2, ALICE studies quark-gluon plasma and the nuclear structure, namely in heavy ion collisions. ATLAS (at IP1) and CMS (at IP5) are general-purpose detectors, with varied physics programs, including the discovery, and subsequent studies of, the Higgs boson. Finally, LHCb, described in this chapter, resides at IP8.

3.2 The LHCb detector in Run 2

The LHCb detector is a single-arm forward spectrometer with a horizontal acceptance of $[10, 300]$ mrad and vertical acceptance of $[10, 250]$ mrad. This gives a coverage of collision products with a pseudorapidity of $2 < \eta < 5$.

The design exploits the natural favouring of $b\bar{b}$ pair production in collisions with a significant difference in the momentum fraction of the colliding partons, *i.e.*, in which the $b\bar{b}$ pair is strongly boosted in the forward/backward directions, as pictured in Fig. 3.2. In collisions with a large difference in parton momentum fraction, the total energy of the colliding particles is much smaller, *i.e.*, closer to the b -quark mass, than the collision energy of the protons. The cross-section of $b\bar{b}$ pair production in 13 TeV pp collisions was measured to be $(144 \pm 1 \pm 21) \mu\text{b}$ [57]; the 6 fb^{-1} integrated luminosity of pp collisions recorded by LHCb in Run 2 correspond to over 860 billion $b\bar{b}$ pairs. This abundance of $b\bar{b}$ pairs makes LHCb an ideal experiment for studying heavy flavour physics.

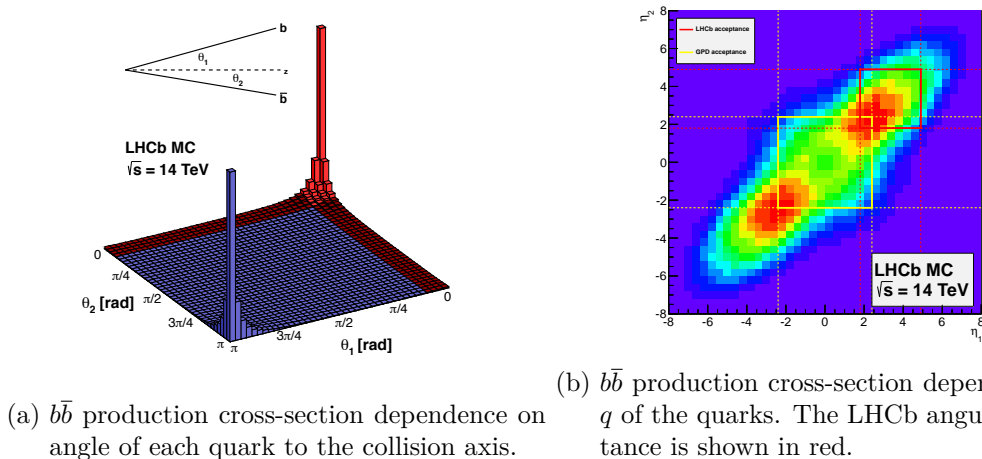


Figure 3.2: $b\bar{b}$ production cross-section dependence on the angle/pseudorapidity of each quark, taken from Ref. [58].

To maintain stable data taking conditions at an approximately constant detector occupancy, the instantaneous luminosity at LHCb is levelled, *i.e.*, the beam dynamics are

varied during each LHC fill to keep the instantaneous luminosity constant [59]. This reduces associated systematic effects and avoids the need to frequently change data-taking conditions such as the trigger configuration (described in Section 3.2.3) by maintaining an approximately constant rate of data from the detector.

As described in Section 2.2.1, according to QCD confinement, the b and \bar{b} quarks cannot be observed as free quarks and instead must hadronise with either valence quarks of the proton or with sea (anti)quarks. The fractions of b -quarks which hadronise with a \bar{u} , \bar{d} , \bar{c} , \bar{s} or ud to form a B^- , \bar{B}^0 , B_c^- , \bar{B}_s^0 or Λ_b^0 , respectively, are known as the fragmentation fractions $f_{u,d,c,s,\text{baryon}}$ [60]. As the sea quarks are produced as $q\bar{q}$ pairs, the \bar{q} must also hadronise; information on the b -hadron flavour can be potentially be inferred from the surrounding hadrons, as described in Section 3.2.3. Since the valence quarks of the protons are always quarks rather than antiquarks, an asymmetry arises in the production of hadrons containing a b/\bar{b} .

The LHCb detector, as operated during Run 2 of the LHC, is shown in Fig. 3.3. The subdetectors can generally be grouped into two categories based on their purpose in the detector: subdetectors responsible for the tracking of charged particles and reconstruction of vertices, discussed in Section 3.2.1, and subdetectors responsible for the identification of particles and calorimetry, discussed in Section 3.2.2. Hardware and software trigger systems are employed to process and reduce the data read out from the detector, which is further processed offline, as described in Section 3.2.3.




Figure 3.3: Side-on schematic of the LHCb detector in Run 2 of the LHC from Ref. [61].

3.2.1 Tracking of charged particles and reconstruction of vertices

Precise measurements of the properties of charged particles tracks and decay vertices are crucial to the LHCb physics programme. A large dipole magnet is employed to deflect any charged particles traversing the detector, such that their momenta can be determined from the measured deflection angle: positively and negatively charged particles are deflected in opposite directions; particles with a smaller momentum are deflected through a larger angle. The Vertex Locator (VELO) and tracker Turicensis (TT) measure trajectories prior to deflection; tracking stations (T1-T3), which lie downstream of the magnet, measure trajectories after deflection. The VELO has the additional task of precisely measuring the decay vertices detached from the collision. Acting together, these achieve measurements of charged particle momenta to precisions below 1% and reconstruct tracks with efficiencies consistently above 95%, as shown in Fig. 3.4. The

3 The LHCb experiment during Run 2 of the LHC

reduced momentum precision at high momenta makes track reconstruction more difficult, as evidenced by the corresponding reduction in track reconstruction in long track reconstruction $p > 100 \text{ GeV}/c$ efficiency in Fig. 3.4b.

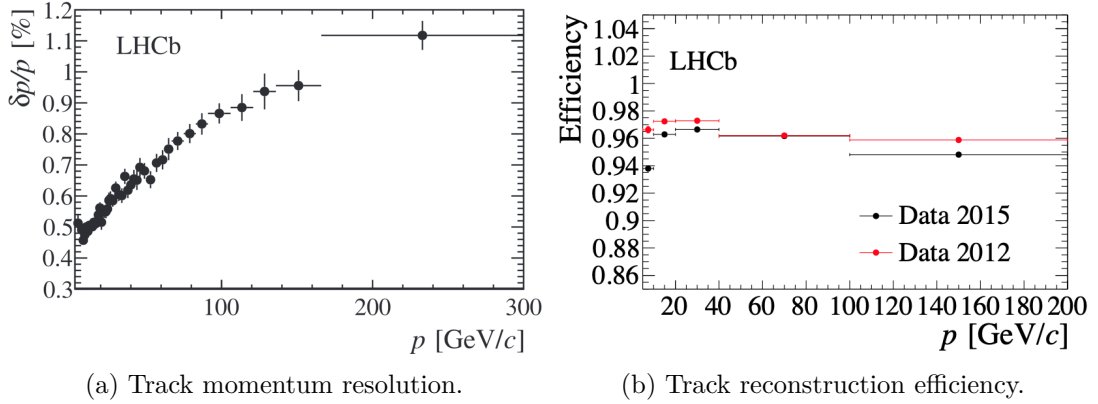


Figure 3.4: Performance of long track (see Fig. 3.5) reconstruction in Run 2, taken from Ref. [62].

Six types of tracks are reconstructed, five of which are shown in Fig. 3.5, with long tracks, *i.e.*, those which contain hits in the VELO and T1-T3 (and optionally TT), being used in most analyses. VELO tracks are straight tracks formed of hits in the VELO, which do not extend into other subdetectors, whilst those extending into the TT are referred to as upstream tracks. Straight tracks solely in T1-T3 are referred to as T tracks, with those which also can be extrapolated to the TT (but not the VELO) dubbed downstream tracks, as these lie downstream of the VELO. Finally, muon tracks, not shown in Fig. 3.5, are simply long tracks which also extend into the muon stations. The track reconstruction algorithm is discussed in more detail in Section 3.2.3

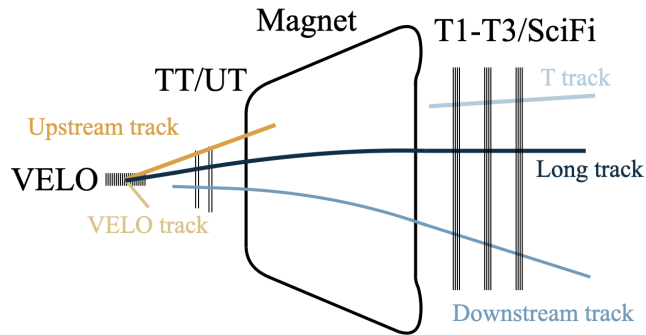


Figure 3.5: Top-down schematic of the track types defined in LHCb adapted from Ref. [61]. The upstream tracker (UT) and scintillating fibre (SciFi) tracker, employed in Run 3, are introduced in Section 5.1.1. Muon tracks, not shown in this schematic, are simply long tracks with hits in M1-M5.

LHCb employs an iron yoke dipole magnet with an integrated field strength of 4 Tm to deflect particles in the horizontal plane [63]. The direction of deflection is periodically reversed to minimise the impact of local asymmetries in the detection of particles. These configurations are labelled as *up* and *down* polarities, based on the direction of the magnetic field.

The VELO is a silicon-strip detector placed around the beam at the IP, with the primary purpose of resolving detached vertices [64]. The VELO is formed of 25 semicircular modules, arranged in layers in the left and right halves of the detector. The halves can be moved toward/away from the beams, such that the modules are held at a safe distance from the beams until the beams are declared stable, and can then be placed as close as 7 mm from the beams (the closest of any sub-detector at the LHC) to improve the resolution of primary vertices (PVs) [65]. This achieves a spatial resolution for clusters of $\sim 4 \mu\text{m}$ nearest the beam (increasing linearly with radial distance) [66]. The consequent IP resolution, on which the selection of displaced vertices strongly depends, is $(15 + 29/p_T) \mu\text{m}$ for p_T in GeV/c [67], where the linear dependence on $1/p_T$ arises from multiple scattering [68]. This is demonstrated, along with the resolution of PV position, in Fig. 3.6. Improvements in these resolutions between Runs 1 and 2 arise predominantly from improvements in the reconstruction algorithm made to enable the reconstruction to take place online, as discussed in Section 3.2.3.

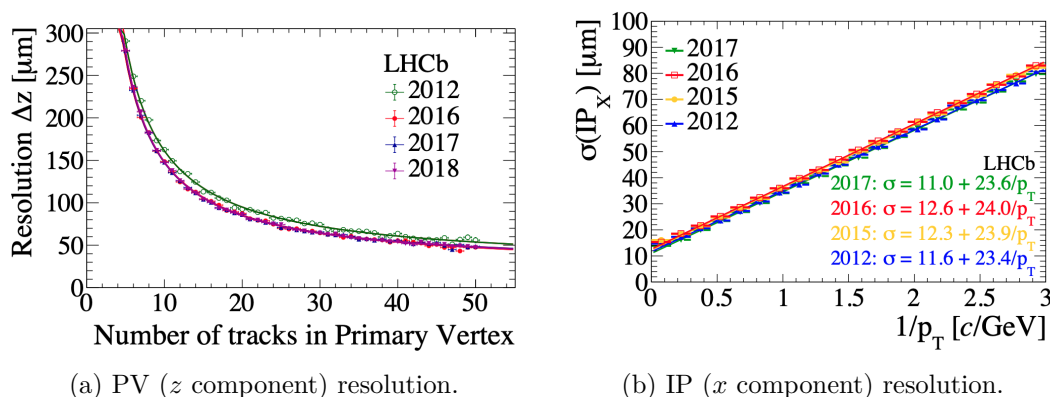


Figure 3.6: Track and vertex reconstruction performance, including comparisons between datasets from Run 2 and the 2012 dataset, taken from Ref. [62].

A beam condition monitor (BCM) [69] is installed to ensure the safety of the subdetectors, in particular the VELO as the closest subdetector to the beam. The BCM consists of two stations, each employing eight symmetrically-placed diamond sensors, situated either side of the IP.

Tracking upstream of the dipole magnet is aided by the TT; tracking downstream of the magnet is handled by T1-T3, each of which consists of an Inner Tracker (IT) [70] and Outer Tracker (OT) [71]. The TT and IT lie close to the beam pipe and thus require a fine granularity due to the high detector occupancy, achieved by their construction of silicon strip detectors, which provide hit resolutions of $\sim 53 \mu\text{m}$, respectively [66]. In contrast, the OT is built as a drift-tube gas detector with a hit resolution of $\sim 200 \mu\text{m}$ [72]. Both the TT, IT and OT all consist of four layers, which are arranged in an X-U-V-X geometry, where X signifies a layer aligned with strips/tubes aligned along the y -axis, *i.e.*, for optimally measuring positions in the x direction, whilst U and V signify layers aligned at $+5^\circ$ and -5° from the y -axis, respectively. The U and V layers provide both the x and y positions, which, when combined with the x positions measured by the X layers and the known separation of the layers, can be used to reconstruct the tracks in all 3 dimensions.

The combined performance of the tracking systems enable the reconstruction of detached b -hadrons with a resolution on their invariant masses down to $8 \text{ MeV}/c^2$.

3.2.2 Particle identification and calorimetry

The LHCb physics programme depends upon high quality identification and calorimetry of final-state particles, namely electrons, muons, charged light hadrons and photons. To facilitate this, several subdetectors are dedicated to particle identification and calorimetry. Two ring image Cherenkov (RICH) detectors measure the angles of Cherenkov radiation emitted by charged particles passing through the detector. Calorimetry of hadrons, electrons and photons is performed by the scintillating pad detector (SPD), preshower (PS), electromagnetic calorimeter (ECAL) and hadronic calorimeter (HCAL). Finally, 5 muon chambers (M1-M5), placed furthest in the detector from the IP.

The RICH detectors are situated either side of the magnet, with RICH1 upstream and RICH2 downstream providing coverage for low- and high-momentum particles, respectively [73]. As a charged particle passes through a material (radiator) of refractive index n , it emits Cherenkov radiation in a cone of angle θ , given by

$$\theta = \cos^{-1} \left(\frac{1}{\beta n} \right) = \cos^{-1} \left(\frac{m}{pn} \right), \quad (3.1)$$

wherein the velocity β is approximated by $\beta = p/m$. Both RICH detectors are filled with gas radiators and employ plane and spherical mirrors to focus the cones of photons onto arrays of hybrid photon detectors (HPDs) [73]. RICH1 uses a radiator of C_4F_{10} , for which $n = 1.0014$, and is optimised for particles with momenta between 2 and 60 GeV/c. RICH2 uses a CF_4 radiator, with $n = 1.0005$, and is optimised for particles with momenta between 15 and 200 GeV/c. The measured Cherenkov angles of muons, pions, kaons and protons on the particle momentum are shown in Fig. 3.7.




Figure 3.7: Cherenkov angle in RICH2 as a function of track momentum, from Ref. [74].

The SPD and PS are scintillating pads separated by a lead absorber, such that electrons and photons begin to form electromagnetic showers in the absorber [75]. Electrons and photons can be distinguished from one another as only electrons produce hits in both the SPD and PS, whilst photons only deposit energy in the PS through the electromagnetic shower. Charged hadrons leave hits in both the SPD and PS, but do not shower in the absorber. A threshold on the maximum number of SPD hits is imposed in L0, such that the busiest $\sim 10\%$ of events are removed, as these events would take disproportionately longer for the HLT to process.

The ECAL and HCAL are both Shashlik calorimeters, with alternating scintillating and absorbing layers, using lead absorbers in the ECAL and iron absorbers in the HCAL [75].

These measure the deposited energy to resolutions of

$$\frac{\sigma_E}{E} = \frac{10\%}{\sqrt{E}} \oplus 1\%, \quad (3.2a)$$

$$\frac{\sigma_E}{E} = \frac{70\%}{\sqrt{E}} \oplus 10\%, \quad (3.2b)$$

respectively, where the first term corresponds to stochastic fluctuations in the particle showers and the second, constant term corresponds to effects such as miscalibrations, non-linearities and leakage [76].

The ECAL can also be used to recover bremsstrahlung photons emitted by electrons as they traverse the detector. Upstream tracks are extrapolated to the ECAL, with any neutral ECAL clusters of $E_T > 75$ MeV which are compatible with the extrapolated track then incorporated in the kinematics of the electron. This improves the momentum resolution of electrons and reduces the long radiative left tail seen in distributions of invariant masses containing an electron, *e.g.*, in Chapter 6. However, this method relies on the bremsstrahlung photons being emitted within the detector acceptance and of a sufficiently high E_T [77, 78].

The muon stations, M1-M5, are the furthest downstream subdetectors. Each of M2-M5 is formed of four concentric Multi Wire Proportional Chamber (MWPC) regions, R1-R4, with increasing granularity towards the innermost region, R1. M1, the only muon station to lie upstream of the calorimeter systems, is constructed similarly; however, since the flux of particles upstream of the calorimeters is much higher, the R1 region is a triple Gas Electron Multiplier instead of an MWPC [79]. 80 cm iron absorbers are placed between each of M2-M5, such that only muons propagate as far as M2-M5. Placing M1 and M2-M5 either side of the magnet allows for the L0 trigger to make selections based on the momenta of muons, as these are deflected by the magnet.

The information from the particle identification subdetectors is combined and evaluated for a given track against the possible particle hypotheses: K^\pm , π^\pm , p , e^\pm and μ^\pm . Likelihoods for each hypothesis are computed, taken as the product of the likelihoods from the RICH, ECAL/HCAL and muon system information individually. PID variables, PID_i are constructed as the difference log-likelihood between each non-pion hypothesis and the pion hypothesis, $\Delta\mathcal{L}(i - \pi)$.

An alternative approach is also taken, in which PID detector information and kinematic/topological information from reconstructed tracks are fed to a set of pretrained Neural Networks (NNs), each of which provides a score on $[0, 1]$ that a particle matches one of the five identities given above or is a ghost, *i.e.*, detector hits reconstructed as tracks but with no real track present [80]. These scores are referred to as the **ProbNN** variables and, in contrast to PID_i , these make direct use of the kinematic and topological features of the tracks of interest.

The performance of the PID and **ProbNN** variables for kaon hypotheses is demonstrated in Fig. 3.8, in particular, their efficiencies and fake rates (the efficiency to select particles which do not match the target hypothesis). In the case on Fig. 3.8, the target hypothesis is a K^\pm , with the efficiency studied in kaons and the fake rate then studied in pions. Both approaches enable efficient selections of each particle hypothesis and allow for clear separations between them in analysis.

3 The LHCb experiment during Run 2 of the LHC

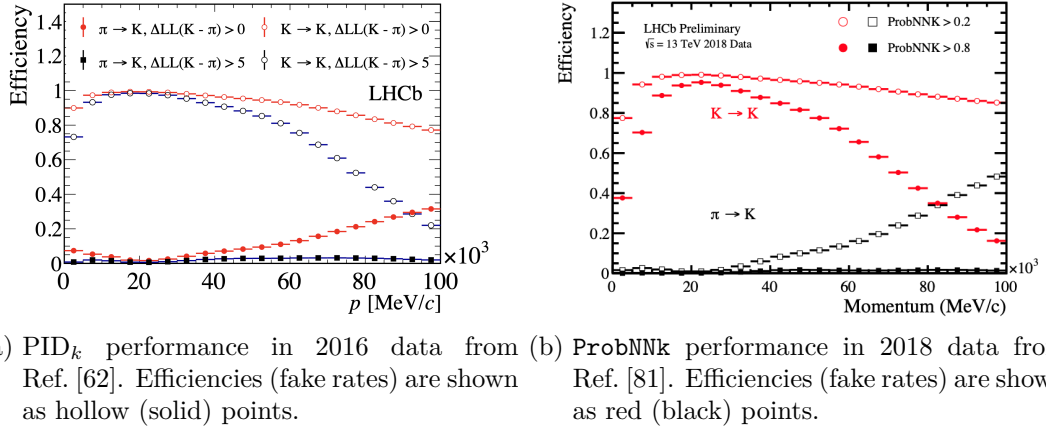


Figure 3.8: Efficiencies and fake rates of PID_k and ProbNNk variables at loose and strict working points, as evaluated during Run 2.

Additional variables `hasRich` and `isMuon` represent whether a particle has corresponding hits in the RICH or muon stations, respectively, whilst `inEcal` represents whether a particle is within the acceptance of the ECAL.

3.2.3 The LHCb data flow

The sheer volume of information produced by the detector readout at the bunch crossing rate of 40 MHz is several orders of magnitude larger than can be feasibly processed and recorded. As is standard practice for LHC experiments, a data-processing system (trigger) is employed to reduce the data rate as efficiently as possible by identifying events of interest, such that the processed data can be stored and further processed offline [82]. Similarly, large quantities of Monte Carlo (MC) simulated LHC collisions must also be processed, stored and distributed. The flow of data, both from the detector and simulation, to the end user performing a physics analysis is shown in Fig. 3.9.

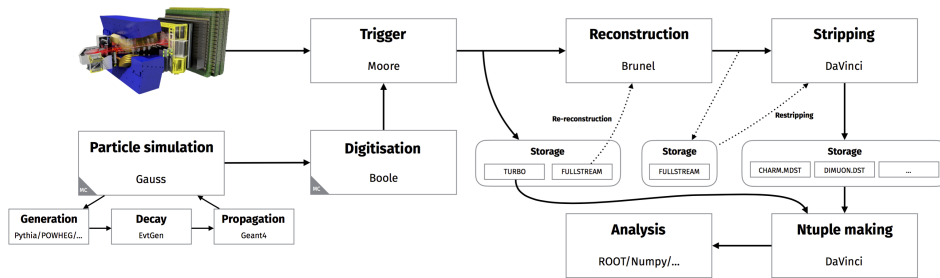


Figure 3.9: Schematic of the LHCb Run 2 dataflow from Ref. [83].

This trigger is split between the Level-0 (L0) hardware trigger and two software triggers, High-Level Trigger 1 (HLT1) and 2 (HLT2) [62], as laid out in Fig. 3.10. L0 must make an initial reduction in rate of $40\times$, with the L0 prompting the full detector to be read out at a more manageable rate of ~ 1 MHz. HLT1 performs a partial reconstruction of events, making decisions based on selection algorithms applied to tracks and simple vertices to provide a further reduction of $\sim 9\times$ to ~ 110 kHz. HLT2 performs a full event reconstruction, enabling more complex selection algorithms to be applied, resulting in a final reduction of $\sim 9\times$ down to ~ 12.5 kHz. Events passing HLT1 are stored temporarily

on a 10 PB disk buffer before being processed by HLT2, with many events being used to align and calibrate the subdetectors in real time

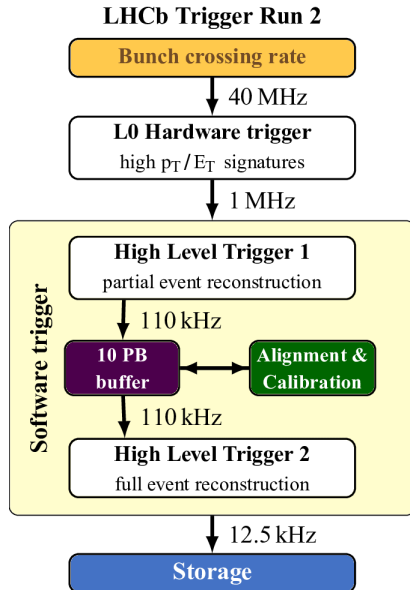


Figure 3.10: Schematic of the LHCb Run 2 trigger systems from Ref. [62].

The L0 hardware trigger must carry out a simple event selection at 40 MHz, relying directly on subdetector information and simple muon tracks to reduce the event rate to 1 MHz. L0 is thus implemented entirely in field-programmable gate arrays, capable of the required rapid decision-making. L0 selects events with high momentum muon tracks or high transverse energy, E_T , calorimeter clusters (2×2 cells in the ECAL or HCAL), as these are typical of heavy flavour decays [84].

The former is used for the L0Muon and LODiMuon selections, whilst the latter is used for the L0Hadron, LOElectron and L0Photon lines. The L0Muon and LODiMuon selections require a single muon track and a pair of related muon tracks, respectively, with the muon tracks required to have a p_T above a given threshold. A distinction between hadrons, electrons and photons is made using information of the SPD and PS, with this distinction defining the L0Hadron, LOElectron and L0Photon selections. Events with a high detector occupancy, which may be very slow to process in HLT1/2, are removed by requiring that events for all lines except LODiMuon have fewer than 450 SPD hits. The efficiencies of the L0 trigger selections in Run 2 in key channels of interest to the LHCb physics programme are shown in Fig. 3.11.

The events selected by L0 are passed directly to the two stages of the HLT, which are operated on the 1700-node Event Filter Farm (EFF) computing farm. The two stages, HLT1 and HLT2, perform a partial and full event reconstruction, respectively, and apply requirements on the reconstructed objects, implemented in the Moore software framework [62]. This enables more granular selections than those applied in L0. These selection algorithms, known as trigger lines, can be either inclusive, targeting a generic signal such as a single muon with high p_T , or exclusive, seeking specific signals such as a $D^0 \rightarrow K^- \pi^+$ decay.

HLT1 receives events from L0 at ~ 1 MHz, reconstructing VELO tracks, which are used to form PVs and are also extrapolated to the remaining tracking stations. Long tracks are

3 The LHCb experiment during Run 2 of the LHC




Figure 3.11: Efficiencies of the L0 trigger selections on key channels of interest to the LHCb physics programme, taken from Ref. [62].

formed by applying pattern recognition algorithms to extrapolate VELO tracks to the TT and T1-T3. These are then fitted using a Kalman filter to improve the momentum resolution and remove some ghost tracks [85]. This information is often sufficient to select high-quality displaced tracks with high p_T , which are likely to originate from b - or c -hadrons. For example, the inclusive HLT1TrackMVA and HLT1TwoTrackMVA lines (labelled HLT1(Two)TrackMVA throughout), which apply multivariate (MVA) classifiers to identify high-quality long tracks displaced from the IP and a pair of such tracks forming a high-quality vertex, respectively. The efficiencies of these lines in Run 2 are shown in Fig. 3.12. Events selected by HLT1 are written to a disk buffer of 10 PB at a rate of ~ 110 kHz.

The events written to the buffer by HLT1 can be used to align and calibrate the subdetectors during data-taking, providing accurate and up-to-date calibration constants for the full reconstruction performed by HLT2. The number of events required to align each subdetector varies, with the alignment and calibration timeframes of a typical LHC fill laid out in Fig. 3.13.

HLT2 performs full reconstruction of events, starting with the reconstruction of charged particle tracks, from which neutral particles are then reconstructed, and finally particle identification is performed. HLT2 runs a mixture of inclusive and exclusive selections. The primary inclusive selections are those of the topological b -hadron trigger, discussed in more detail in Section 5.2.1, which apply MVAs to the kinematic and topological properties of b -hadron decay candidates, the efficiencies of which are shown in Fig. 3.14.

The detailed selections enabled by the full reconstruction allow HLT2 to reduce the event rate to ~ 12.5 kHz, corresponding to a bandwidth of ~ 0.6 GB/s. To record more events within the same bandwidth, the Turbo stream was introduced in Run 2, wherein only the reconstructed signal candidate of a given event is recorded, rather than the entire event; events recorded in full are instead saved to the Full stream. This approach, commonly referred to as selective persistency, wherein events are recorded with reduced event content, is schematised in Fig. 3.15. In Ref. [62], the average size of a full event was 48 kB, whilst the average size for a Turbo event was only 7 kB.

Samples collected for calibration, such as those of the PID/trigger response and track

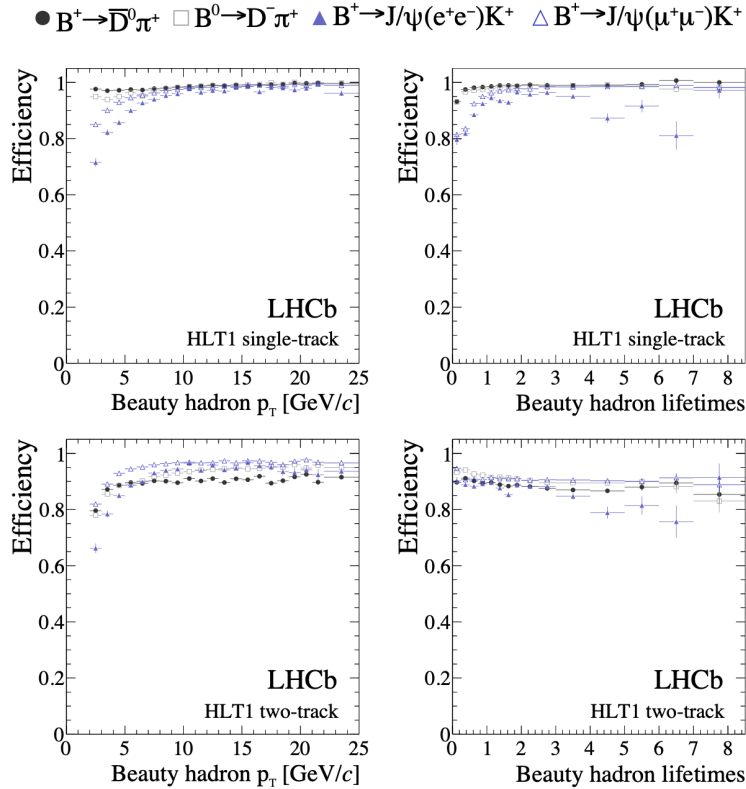


Figure 3.12: Efficiencies of HLT1TrackMVA and HLT1TwoTrackMVA on key channels of interest to the LHCb physics programme, taken from Ref. [62].

reconstruction used in Section 4.2, are collected using the Turbo stream [86]. Selection algorithms intended to collect candidates for direct use in physics analyses were also implemented for use with the Turbo stream, though these contributed only a small fraction of the HLT2 output bandwidth.

Events saved to the Full stream are subject to an additional event reconstruction with the `Brunel` software framework and further application of selections with the `DaVinci` framework [88], saving the resulting data into `Root` NTuples [89]. The application of these additional selections, dubbed stripping, which are near to analysis-quality and typically target a single signal or small group of related signals, reduce the sizes of samples provided for physics analysis [90]. Additional detailed information about the event can be included at this stage, *e.g.*, the projections of particle positions into the HCAL as discussed in Section 4.2.3. An important input commonly included at this stage is the information on surrounding particles in the event. As described in Section 3.1, these particles may be sufficient to infer the flavour of a b -hadron on production, necessary for studies where the original hadron flavour must be known, such as time-dependent CP violation measurements, *e.g.*, in Refs. [32, 91]. This inference, known as flavour tagging, applies pre-trained NNs to identify either kaons, protons or pions in the vicinity of the b -hadron of interest or another b -hadron decay in the event consistent with that of the hadron formed from the corresponding \bar{b} [92]. As this relies on reconstructing and successfully identifying the additional particles in the event, requiring tagging information involves a notable reduction in statistics.

Monte Carlo (MC) simulations of LHC collisions are used extensively in both analysis and detector development. Simulated samples contain truth-level information on the

3 The LHCb experiment during Run 2 of the LHC




Figure 3.13: Alignment and calibration procedure for a typical LHC fill in Run 2, indicating the points at which sufficient data has been accumulated for each process, taken from Ref. [62].




Figure 3.14: Efficiencies of the topological b -hadron trigger in HLT2 on key channels of interest to the LHCb physics programme, taken from Ref. [62].

properties of each event, allowing for studies comparing the effect of reconstruction and selection in channels of interest, *e.g.*, in Sections 4.2.1 & 6.3. These studies often include evaluation of differences between MC simulated and data samples, as complex properties such as PID or trigger responses are non-trivially related to the fundamental properties of the event and thus non-trivially related to effects from mismodelling. Additionally, distributions of variables in MC samples can be used to provide templates to fits in analyses, *e.g.*, in Sections 4.3.1 & 6.2.

The simulation of events in the **Gauss** framework is discussed in detail in Ref. [93] and summarised briefly here. **Pythia** [94,95] is used to simulate pp collisions with an LHCb-specific configuration [96], with the hadronisation and decay of, and radiation of photons from the resulting particles propagated by **EvtGen** [97] and **Photos** [98]. The interaction of these particles with the LHCb detector and subsequent detector response is simulated by **Geant4** [99]. **Boole** is then used to digitise the detector response, including the L0 trigger. As in Fig. 3.9, MC simulated samples are subject to the same data processing as data read out from the detector. Since the true identities of simulated candidates are known, the reconstructed simulated events can be compared to their true identities and assigned a corresponding background category [100]. By making requirements (known as

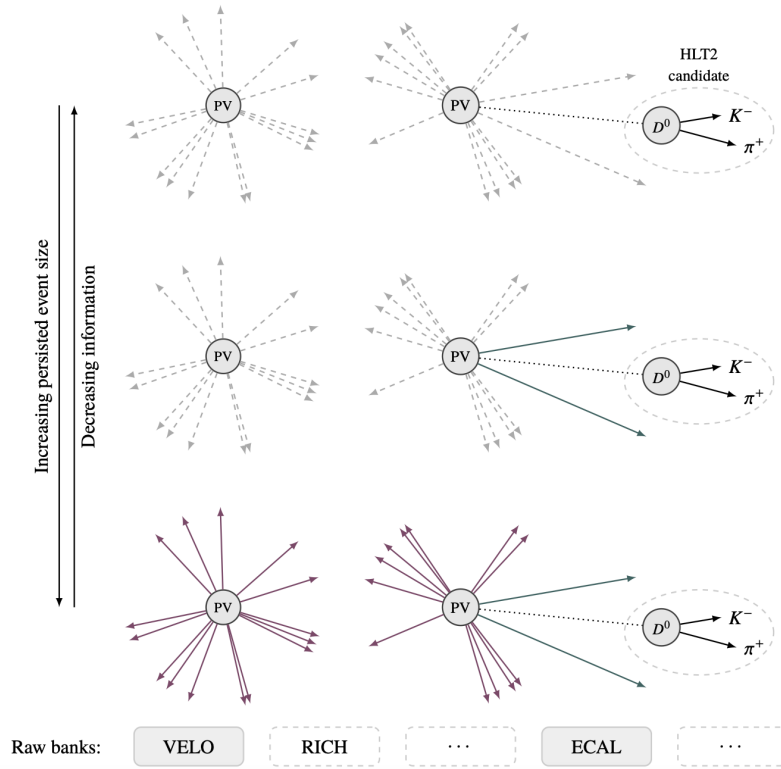


Figure 3.15: Schematic of the Turbo event model, in which objects can be persisted as needed, ranging from only the HLT2 candidate (top) to the full event (bottom), taken from Ref. [87].

truth-matching) on this background category, MC simulated candidates can be filtered, *e.g.*, to select all candidates which are correctly reconstructed signal, or those which are constructed missing a particle from the true decay chain.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

Measurements of CP asymmetries can probe the existence of NP indirectly, as laid out in Section 2.3. This chapter covers precisely such a test of the SM, measuring the decay-time-integrated flavour-untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays, from which a direct CP violating contribution to A_{fs}^s (per Eq. 2.47) can be extracted. [37]. The presence of such a contribution would indicate whether the $b \rightarrow c\bar{u}q$ anomaly, discussed in Section 2.4.1, is caused by NP.

This measurement, which is performed using 5.4fb^{-1} of pp collision data recorded by LHCb between 2016 and 2018, requires the asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays in the data to be extracted and corrected for contributions from production and detection asymmetries. To extract the asymmetry, the nature of contributions from $B_s^0 \rightarrow D_s^- \pi^+$ decays and a collection of background sources in the data must be understood in detail. Similarly, determination of the production and detection asymmetries require a deep understanding of the collider and detector effects at play.

The author is a proponent and contact author of this measurement (which is in internal review within LHCb at time of submission), and has made significant contributions throughout the work discussed in this chapter.

Note: since the submission of this work, the measurement described in this chapter has been made public. For more details see Appendix A.1.

4.1 Analysis strategy, samples and selections

The CP asymmetry A_{fs}^s , as given in Eq. 2.47 in Section 2.2.4, is time-dependent, varying with $B_s^0\text{-}\bar{B}_s^0$ oscillation. As such, measuring A_{fs}^s would require knowledge of the initial flavour of the meson, which can be obtained by means of flavour tagging (per Section 3.2.3). However, since the effective tagging efficiencies for $B_{(s)}^0$ at LHCb are $\mathcal{O}(1\%)$ [101, 29], such an observable is not experimentally favourable. Instead, the flavour-untagged asymmetry, A_{untagged}^s , which takes the form of Eq. 2.36, is considered.

In the SM, $B_s^0 \rightarrow D_s^- \pi^+$ decays proceed only via the tree-level $b \rightarrow c\bar{u}d$ transition. The effect of an additional amplitude from a hypothetical NP process can be determined by taking the amplitudes from Eqs. 2.41a & 2.41b and their \bar{B}_s^0 counterparts, removing terms containing $\bar{\mathcal{A}}_f$ and $\mathcal{A}_{\bar{f}}$ as $B_s^0 \rightarrow D_s^- \pi^+$ is flavour-specific [102]. Noting from Eq. 2.40 that $|q/p| \approx 1 - a_{\text{fs}}^q/2$, this can be derived as

$$A_{\text{untagged}}^s \approx \frac{2r \sin(\phi) \sin(\delta) - a_{\text{fs}}^q (1 + 2r \cos(\phi) \cos(\delta)) Y(t)}{1 + 2r \cos(\phi) \cos(\delta) - 2a_{\text{fs}}^q r \sin(\phi) \sin(\delta) Y(t)}, \quad (4.1)$$

wherein $Y(t)$ is given by

$$Y(t) = \frac{f_-(t)}{f_+(t) + f_-(t)} = \frac{1}{2} \left(1 - \frac{\cos(\Delta m_q t)}{\cosh\left(\frac{\Delta\Gamma_q t}{2}\right)} \right). \quad (4.2)$$

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

Measuring the decay-time-integrated equivalent, $\langle A_{\text{untagged}}^s \rangle$, this integrated over t to give

$$\begin{aligned} \langle A_{\text{untagged}}^s \rangle &\approx \frac{4r \sin(\phi) \sin(\delta) - a_{\text{fs}}^q (1 - \rho_q) (1 + 2r \cos(\phi) \cos(\delta))}{2(1 + 2r \cos(\phi) \cos(\delta) - a_{\text{fs}}^q (1 - \rho_q) r \sin(\phi) \sin(\delta))} \\ &\approx A_{\text{dir}}^q - \frac{a_{\text{fs}}^q}{2} (1 - \rho_q), \end{aligned} \quad (4.3)$$

where we have used the fact that both r and a_{fs}^q are small, and wherein ρ_s describes the relative oscillation frequency,

$$\rho_q = \int Y(t) dt = \frac{\Gamma_q^2 - \frac{\Delta\Gamma_s}{4}}{\Gamma_q^2 + \Delta m_s^2} \approx 0.001. \quad (4.4)$$

On this basis, the sensitivity to A_{fs}^s is reduced by a factor of 1/2, though still significantly greater if flavour tagging were to be applied. From $\langle A_{\text{untagged}}^s \rangle$, it is therefore possible to infer the presence of a general NP contribution in a model-independent way, as encapsulated in A_{dir}^s . Given the precision of current constraints on a_{fs}^s , a measurement of $\langle A_{\text{untagged}}^s \rangle$ of $\mathcal{O}(10^{-2})$ would provide evidence for a BSM explanation of the $b \rightarrow c\bar{u}q$ anomaly [37].

As in Ref. [37], $\langle A_{\text{untagged}}^s \rangle$ can be expressed as a sum of asymmetry terms:

$$\langle A_{\text{untagged}}^s \rangle = A_{\text{raw}} - A_{\text{det.}} - \underbrace{A_{\text{prod.}}}_{A'_{\text{prod.}}} = A_{\text{raw}} - A_{\text{det.}} - \frac{\int e^{-\Gamma_s t} \cos(\Delta m_s t) \varepsilon(t) dt}{\int e^{-\Gamma_s t} \cosh\left(\frac{\Delta\Gamma_s t}{2}\right) \varepsilon(t) dt}, \quad (4.5)$$

where A_{raw} , $A_{\text{det.}}$ and $A_{\text{prod.}}$ are the raw, detection and production asymmetries.

The raw CP asymmetry is measured between $B_s^0 \rightarrow D_s^- \pi^{+1}$ decays to $D_s^+ \pi^-$ and $D_s^- \pi^+$, defined as

$$A_{\text{raw}} = \frac{N(D_s^+ \pi^-) - N(D_s^- \pi^+)}{N(D_s^+ \pi^-) + N(D_s^- \pi^+)}, \quad (4.6)$$

for observed yields $N(D_s^\pm \pi^\mp)$. This can be extracted from the data by means of a negative-log likelihood (NLL) fit to the $hh\pi\pi$ invariant mass, discussed in Section 4.3.

The inherent CP asymmetry in detecting the final-state particles, $A_{\text{det.}}$, can be estimated by considering the individual asymmetries effects in the detection, reconstruction and selection of final-state particles, studying these in MC simulated samples, using reference samples and calibration tables to determine these, discussed in Section 4.4. As the final state is charge symmetric, contributions to the detection asymmetry should largely cancel, especially in the h^+h^- pair. The π^+ of the B_s^0 and π^- of the D_s^- have significantly different kinematics and so such effects are unlikely to cancel completely.

The inherent CP asymmetry between the production of the B_s^0 and \bar{B}_s^0 in pp collisions, $A_{\text{prod.}}$, discussed in Section 3.2, has previously been measured in $B_s^0 \rightarrow D_s^- \pi^+$ decays from $\sqrt{s} = 13$ TeV pp collisions as $(-0.31 \pm 0.32)\%$ in Ref. [29], though this is washed out further by the rapid $B_s^0 - \bar{B}_s^0$ oscillations, quantified by a factor which can be numerically computed from the decay time acceptance function (also given in Ref. [29]), discussed in Section 4.5.

¹Charge conjugation of decays is assumed henceforth, except where explicitly stated otherwise.

4.1.1 Data samples and trigger selections

The measurement of $\langle A_{\text{untagged}}^s \rangle$ described here uses the dataset of pp collisions recorded by LHCb between 2016 and 2018, corresponding to an integrated luminosity of 5.4 fb^{-1} . The integrated luminosities of the constituent years of data-taking in each magnet configuration are listed in Table A.1 in Appendix A. The large sample of $B_s^0 \rightarrow D_s^- \pi^+$ decays collected during Run 2 was notably used to achieve the highest precision measurement yet of the $B_s^0 - \bar{B}_s^0$ difference, Δm_s [29].

Firstly, L0Hadron is required to trigger on (TOS²) any of the final-state particles in each candidate. At HLT1, either of the lines HLT1TrackMVA or HLT1TwoTrackMVA, is required to trigger on the B_s^0 candidate. No HLT2 requirement is applied offline; however, the stripping lines B02D{K,Pi}D2HHHBeauty2CharmLine used to produce the samples require that either HLT2TopoTwoBody or HLT2TopoThreeBody triggers on the B_s^0 candidate. The decay is reconstructed by first combining three long tracks to form a good quality vertex (that of the D_s^+), and then attaching another long track with a π^\pm mass hypothesis to form another vertex (that of the B_s^0) which must be displaced from the PV.

Both of the dominant hadronic decay modes of the D_s^- are studied in the analysis: $D_s^- \rightarrow K^- K^+ \pi^-$ and $D_s^- \rightarrow \pi^- \pi^+ \pi^-$, with the measurement performed independently in each mode. The tracks can therefore be assigned either the $KK\pi\pi$ or $\pi\pi\pi\pi$ mass hypothesis. As a convention, the π^+ produced in the B_s^0 decay (the π^+ in $h^- h^+ \pi^- \pi^+$) is referred to throughout as the associated pion (labelled $\pi_{B_s^0}$), whilst the hadrons produced in the $D_s^- \rightarrow h^+ h^- \pi^+$ decays are referred to as the (D_s^-) daughter hadrons.

No distinction is made on how each $D_s^- \rightarrow h^- h^+ \pi^-$ proceeds: in $D_s^- \rightarrow K^- K^+ \pi^-$, the primary contributions are the decays via $\phi \rightarrow K^+ K^-$ and $K^{*0} \rightarrow K^- \pi^+$, and the non-resonant mode; in $D_s^- \rightarrow \pi^- \pi^+ \pi^-$, these are decays via $f_0(980) \rightarrow \pi^+ \pi^-$ and the non-resonant mode. Windows in the invariant masses of opposite-sign pairs of D_s^- daughter hadrons are used throughout the analysis. The windows associated with the ϕ and K^{*0} are taken to be $m(K^+ K^-) \in [980, 1080] \text{ MeV}/c^2$ and $m(K^+ \pi^-) \in [825, 950] \text{ MeV}/c^2$, respectively. As two opposite-sign $\pi^+ \pi^-$ pairs can be formed in the $D_s^- \rightarrow \pi^- \pi^+ \pi^-$ mode, a meaningful distinction between the pairs can be made by taking the pair with the smaller invariant mass as $m(\pi^+ \pi^-)_{\text{low}}$ and the pair with the larger invariant mass as $m(\pi^+ \pi^-)_{\text{high}}$. As the $f_0(980)$ can produce either $\pi^+ \pi^-$ pair, the corresponding window is taken as $m(\pi^+ \pi^-)_{\text{low,high}} = [930, 1010] \text{ MeV}/c^2$. The square invariant masses of the two $h^+ h^-$ pairs of each D_s^- decay can be plotted against one another to visualise the phase-space of these three-body decays [103], as is shown later in Fig. 4.1.

MC simulated samples of $B_s^0 \rightarrow D_s^- \pi^+$ decays and the background contributions described in the following section were produced for this analysis. These samples are subject to the same online data processing and stripping as the data samples, in addition to truth-matching requirements which ensure that the $B_s^0 \rightarrow D_s^- \pi^+$ samples contain only signal and that the background samples contain their respective type of background, as laid out in the next section.

4.1.2 Background contributions

In reconstructing and selecting $B_s^0 \rightarrow D_s^- \pi^+$ decays, candidates may arise from sources other than the decays of interest. For example, the D_s^- or B_s^0 composites may be

²For a full discussion of how this is defined, see the discussion of the TISTOS method in Section 5.2.2.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

constructed using particles originating from the pp collision or from other decays (combinatorial backgrounds), the identity of final-state particles may be incorrectly assigned (misidentified background), or the D_s^- or B_s^0 composites may be constructed from decays where one or more particle is missing in the reconstruction (partially reconstructed background). The background contributions affecting this analysis, namely those which result in candidates with a $hh\pi\pi$ invariant mass in the vicinity of the B_s^0 mass, are summarised in Table 4.1 and discussed in more detail below.

Table 4.1: Summary of background contributions which may enter the $hh\pi\pi$ invariant mass distributions, and the sources by which they arise. Contributions marked * and † are proven not present in (see Section 4.1.4) and removed entirely (see Section 4.1.3) from the data samples, respectively.

Background	Category	Source
Multiple candidates [†]	Reconstruction	Alternative valid reconstructions
Combinatorial	Combinatorial	Random $hh\pi\pi$ not originating from a B_s^0 or D_s^-
Charmless*		Correct $hh\pi\pi$, but $hh\pi$ not originating from a D_s^-
Track-swap*	Misidentified composite	D_s^- mistakenly constructed using the real associated π^+
$B^0 \rightarrow D_s^+ \pi^-$	particle	Correct $hh\pi\pi$, proceeding via D_s^- , but coming from B^0 instead of B_s^0
$B^0 \rightarrow D^{*-} \pi^+$		Correct (or misidentified) $hh\pi\pi$, from incorrect c -hadron
$B_s^0 \rightarrow D_s^- K^+$		Associated K^+ misidentified as π^+
$B^0 \rightarrow D^- \pi^+$	Misidentified final-state particle	Misidentification of π^+ as K^+ in $D_s^- \rightarrow K^- K^+ \pi^-$ or K^- as π^- in $D_s^- \rightarrow \pi^- \pi^+ \pi^-$
$\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$		Misidentified p as K^+ in $\bar{\Lambda}_c^- \rightarrow \bar{p} K^+ \pi^-$
$B_s^0 \rightarrow D_s^{*-} \pi^+$	Partially reconstructed	Missing π^0 in $D_s^{*-} \rightarrow D_s^- \pi^0$ or γ in $D_s^{*-} \rightarrow D_s^- \gamma$
$B_s^0 \rightarrow D_s^- \rho^{+\dagger}$		Missing π^0 in $\rho^+ \rightarrow \pi^+ \pi^0$

Firstly, from the four (three) long tracks combined to form the B_s^0 (D_s^+) composites, there can often be several valid reconstructions within the same event, especially when considering that there may be multiple ways to assign the $hh\pi\pi$ ($hh\pi$) mass hypothesis, *e.g.*, whether the same-sign long tracks of the D_s^+ candidate are assigned as K^+ or π^+ , or vice versa, for the $KK\pi$ hypothesis. Therefore, it is entirely possible that the same candidate may be reconstructed more than once, and that these candidates may pass the online and offline selections. This background is primarily suppressed by imposing requirements on PID and ProbNN variables, though truly removing this background involves ensuring that each event number in the sample corresponds to only one candidate.

Combinatorial backgrounds are caused by the combination of particles from other sources,

either in constructing the D_s^- or the B_s^0 as a whole. No peaking structures are anticipated from combinatorial backgrounds in the $hh\pi\pi$ invariant mass. A D_s^- mass peak arises in the $hh\pi$ invariant mass distribution for combinatorial background events consisting of a real $D_s^- \rightarrow h^- h^+ \pi^-$ decay with a random π^+ attached, whilst no such peak arises if the tracks forming the D_s^- candidate do not originate from the same real D_s^- .

Charmless background arises from the correct reconstruction of decays of B_s^0 to $hh\pi\pi$ without proceeding via a c -hadron resonance, which then appears as a peaking structure in the $hh\pi\pi$ invariant mass, whilst appearing combinatorial-like in the $hh\pi$ invariant mass. Likewise, $B^0 \rightarrow D_s^+ \pi^-$ can contribute, as these decays are reconstructed correctly and proceed via D_s^- ; however, the original b -hadron is incorrect, resulting in a peak of the same shape as the signal peak, but centred around the B^0 mass rather than the B_s^0 mass. The $hh\pi\pi$ can also be combined incorrectly, such that the associated π^+ is mistakenly used to form the D_s^- and a π^+ from the D_s^- is assigned as the associated pion, though this typically only arises if the mistaken pions are kinematically similar. Similarly, the decay $B^0 \rightarrow D^{*-} \pi^+$ (wherein $D^{*-} \rightarrow D^0 \pi^+$ with $D^0 \rightarrow K^- K^+$ or $D^0 \rightarrow \pi^- \pi^+$) produces the correct final-state particles, albeit proceeding from the wrong b -hadron and via the wrong c -hadron.

The misidentification of final-state hadrons introduces several possible background contributions, wherein decay channels containing other final-state particles may enter. Of particular interest to this analysis are the misidentification of the associated K^- in $B_s^0 \rightarrow D_s^- K^+$ as a π^+ , the misidentification of the π^+ in the D^- decay of $B^0 \rightarrow D^- \pi^+$ as a K^+ , and the misidentification of the p in the Λ_c^+ decay of $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ as a K^+ . These contributions appear as skewed distributions in the $hh\pi\pi$ invariant mass, shifted below or above the mass of their b -hadron depending on the misidentification involved, *e.g.*, $B_s^0 \rightarrow D_s^- K^+$ is shifted below the B_s^0 mass as the π^+ assigned in place of the K^+ has a lower mass. As these shifts are typically below $\mathcal{O}(100 \text{ MeV})$, the contributions overlap significantly with $B_s^0 \rightarrow D_s^- \pi^+$ in the $hh\pi\pi$ invariant mass spectrum, and can thus be difficult to disentangle. $B^0 \rightarrow D^{*-} \pi^+$ can also enter again similarly, as the D^0 decays as $D^0 \rightarrow K^- \pi^+$ (more abundantly than as $D^0 \rightarrow K^- K^+$ or $D^0 \rightarrow \pi^- \pi^+$) and the misidentification of either the π^+ as a K^+ or K^- as a π^- for the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively. Misidentification of e^\pm as h^\pm and misidentification of multiple final state hadrons are not considered in this analysis as these contributions can be suppressed to a negligible level by requirements on particle identification information. Similarly, misidentification of μ^\pm as h^\pm is not considered as requirements on particle identification information and the response of the muon system can efficiently exclude such contributions.

Finally, since the final state of interest is $hh\pi\pi$, no neutral particles are included in the construction of candidates. Thus, decays containing the correct final-state hadrons but also containing additional neutral particles, such as $B_s^0 \rightarrow D_s^{*-} \pi^+$ (wherein $D_s^{*-} \rightarrow D_s^- \gamma$ or $D_s^{*-} \rightarrow D_s^- \pi^0$) and $B_s^0 \rightarrow D_s^- \rho^+$ (wherein $\rho^+ \rightarrow \pi^+ \pi^0$), are reconstructed without the full decay products. These partially reconstructed backgrounds typically appear as peaking structures below the B_s^0 mass in the $hh\pi\pi$ invariant mass distribution, as they lack the additional energy of the missing particle(s).

Whilst the selections enforced in the trigger suppress background significantly, further selections must be applied offline to improve the signal purity of samples. Requirements on the kinematics and topology of the B_s^0 and D_s^- candidates can suppress combinatorial and charmless backgrounds, whilst requirements on PID or `ProbNN` variables of the final-state hadrons may suppress misidentified backgrounds. Suppressing partially

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

reconstructed background efficiently is usually non-trivial, though windows in the $hh\pi\pi$ and $hh\pi$ invariant masses can often reduce or exclude these.

Requirements targeting specific backgrounds, commonly referred to as vetoes, can also be imposed. These are applied when a background component can no longer be suppressed further by a global selection without significant reduction of the selection efficiency on signal.

4.1.3 Offline selection

To mitigate the backgrounds discussed above, a cut-based selection is applied offline to the candidates passing the trigger selections and constructed by the stripping lines. These selection requirements are summarised in Table 4.2 and discussed in detail below. The $hh\pi\pi$ and $hh\pi$ invariant masses, and a Dalitz plot of the $hh\pi$ system of the selected events are shown in Fig. 4.1. Unlike many measurements in $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decay channels, no MVA-based selection is applied here to remove combinatorial background, as the requirements made in Table 4.2 are sufficient and quantifying the effect of such a classifier on A_{raw} would be non-trivial.

Table 4.2: Summary of selection requirements.

Final state	Selection	Requirement	
Both	Trigger	L0Hadron TOS on any h^\pm	
		HLT1(Two)TrackMVA TOS on B_s^0	
		HLT2Topo{Two,Three}Body TOS on B_s^0	
	Stripping	B02D{K,Pi}D2HHHBeauty2CharmLine	
	Unique candidates	Duplicate candidates per event discarded	
	Mass windows	$m(hh\pi\pi) \in [5280, 6000]$ MeV	
		$m(hh\pi) \in [1939, 1999]$ MeV	
	Particle identity	ProbNNk > 0.6 ProbNNpi < 0.4	for K^\pm
		ProbNNk < 0.4 ProbNNpi > 0.6	for $\pi_{B_s^0, D_s^-}^\pm$
		hasRich = 1 isMuon = 0	for h^\pm
HCAL separation	$d_{\text{HCAL}} > 1000$ mm	for all hh pairs	
$KK\pi\pi$	Particle identity in non-resonant phase-space	ProbNNk > 0.8 ProbNNpi < 0.2	for K^\pm
		ProbNNpi > 0.8 ProbNNk < 0.2	for $\pi_{D_s^-}^\pm$
$B^0 \rightarrow D^{*-} \pi^+$ veto		$m(K^+ K^-) \notin [1830, 1890]$ MeV	
		$m(K^- \pi_{B_s^0}^+) \notin [1834, 1894]$ MeV	
$\pi\pi\pi\pi$	Particle identity in non-resonant phase-space	$m(\pi^- \pi_{B_s^0}^+ [\text{as } K^+]) \notin [1834, 1894]$ MeV	
		$m(\pi^- \pi^+)_{\text{high.}} \notin [1834, 1894]$ MeV	
		ProbNNpi > 0.8 ProbNNk < 0.2	for $\pi_{D_s^-}^\pm$

4.1 Analysis strategy, samples and selections

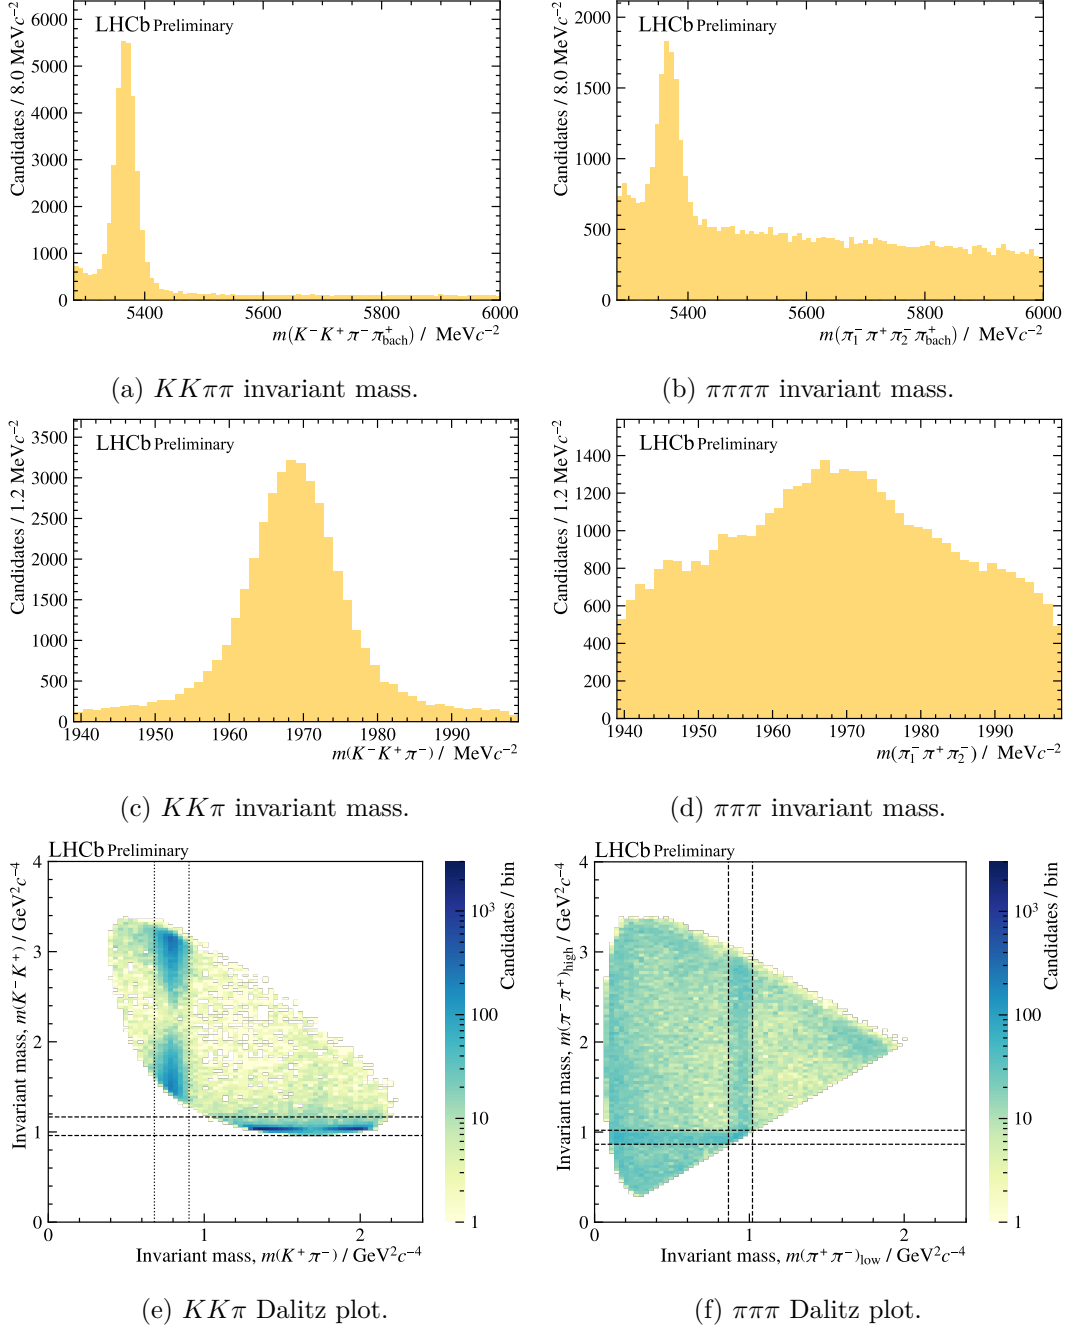


Figure 4.1: Distributions of $hh\pi\pi$ and $hh\pi$ invariant masses, and Dalitz plots of $hh\pi$ phase space in the 2016-18 dataset, following the application of the selection requirements listed in Table 4.2. The $KK\pi\pi$ ($\pi\pi\pi\pi$) final state is shown on the left (right). Windows around the ϕ and K^{*0} in the $D_s^- \rightarrow K^- K^+ \pi^-$ mode and $f_0(980)$ in the $D_s^- \rightarrow \pi^- \pi^+ \pi^-$ mode described in Section 4.1.1 are overlaid as dashed lines.

Firstly, each event is required to correspond to only one candidate. Of the events selected by the requirements discussed below, $(0.474 \pm 0.033)\%$ and $(2.421 \pm 0.068)\%$ of events correspond to more than one candidate in the $KK\pi\pi$ and $\pi\pi\pi\pi$ datasets, respectively. These must be removed, else events will, in effect, be counted twice. For candidates with the same event number within the same sample, all but one of the candidates per event number are removed at random. A similar requirement can also be imposed to ensure that the samples of the two final states are mutually exclusive: any candidates in the $\pi\pi\pi\pi$ sample with event numbers appearing in the $KK\pi\pi$ sample are removed, as $D_s^+ \rightarrow K^+ K^- \pi^+$ has a greater branching fraction.

Firstly, the reconstructed candidates are required to lie within invariant mass windows of $m(hh\pi\pi) \in [5280, 6000]$ MeV and $m(hh\pi) \in [1939, 1999]$ MeV, *i.e.*, the vicinity of the B_s^0 and D_s^+ invariant masses respectively. The former is chosen, to have a lower limit excluding the $B_s^0 \rightarrow D_s^- \rho^+$ background and an upper limit providing sufficient combinatorial background to perform the invariant mass fit described in Section 4.3.2 stably. The latter is taken as $30 \text{ MeV}/c^2$ either side of $m_{D_s^+}$ (reported as $(1968.35 \pm 0.07) \text{ MeV}/c^2$ in Ref. [5]), as the D_s^- peak in the $hh\pi$ invariant mass distribution has a resolution of approximately $\sim 7 \text{ MeV}/c^2$ (see Fig. 4.1). This rejects candidates with an $hh\pi$ invariant mass which is unphysically high or low.


Further to the loose requirements on particle identity metrics made at the HLT2, additional requirements are made to the `ProbNNpi` and `ProbNNk` scores (see Section 3.2.2) of each final-state particle. All final-state π^\pm (K^\pm) are required to have scores of `ProbNNpi` (`ProbNNk`) > 0.6 and `ProbNNk` (`ProbNNpi`) < 0.4 . As the signal purity is inherently lower in the non-resonant regions of the $D_s^+ \rightarrow h^+ h^- \pi^+$ phase space, stricter requirements of `ProbNNpi` (`ProbNNk`) > 0.8 and `ProbNNk` (`ProbNNpi`) < 0.2 are applied to each π^\pm (K^\pm) in these regions, *i.e.*, outside the resonance windows described in Section 4.1.1. Each final-state particle is also required to have a corresponding response in the RICH subdetectors and no response in the muon systems. All possible pairs of final-state particles are required to be separated by at least 1000 mm in the HCAL, or for either/both of the hadrons to lie outside the HCAL acceptance, to enable accurate estimation of the L0Hadron efficiency, discussed in more detail in Section 4.2.

The selections discussed thus far generally provide adequate suppression of most backgrounds; however, vetoes are still considered to mitigate some backgrounds such as that of $B^0 \rightarrow D^{*-} \pi^+$.


Considering the backgrounds listed in Table 4.1, the $hh\pi\pi$ and $hh\pi$ invariant mass windows are already sufficient to entirely exclude the $B^0 \rightarrow D^- \pi^+$ and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ backgrounds in the $\pi\pi\pi\pi$ final state, and the $B_s^0 \rightarrow D_s^- \rho^+$ background, as demonstrated in Fig. 4.2. Examining MC simulated samples of $B_s^0 \rightarrow D_s^- \rho^+$ decays, only $(0.196 \pm 0.010)\%$ of candidates fall inside the $hh\pi\pi$ window. These windows also remove the more complex structures of the $B_s^0 \rightarrow D_s^{*-} \pi^+$ component in $m(hh\pi\pi)$, with only the furthest right tail (which, as demonstrated later in Section 4.3, can be modelled by a standard peaking structure) falling within the window. In MC simulated $B_s^0 \rightarrow D_s^{*-} \pi^+$ decays, $(4.396 \pm 0.045)\%$ of candidates lie within the $hh\pi\pi$ and $hh\pi$ windows.

Whilst vetoes on $B^0 \rightarrow D^- \pi^+$ and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ were considered in the $KK\pi\pi$ final state, *e.g.*, applying stricter `ProbNN` requirements for candidates with a $KK\pi$ invariant mass lying near to the Λ_c^+ mass when the K^+ is assigned the p mass, these backgrounds can be modelled accurately from MC simulated samples and thus were not suppressed further. Not employing vetoes against these background processes ensures that sufficiently many


4.1 Analysis strategy, samples and selections




(a) $hh\pi\pi$ for $B_s^0 \rightarrow D_s^- \rho^+$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$.




(b) $hh\pi\pi$ for $B^0 \rightarrow D^- \pi^+$.



(c) $hh\pi$ for $B^0 \rightarrow D^- \pi^+$.



(d) $hh\pi\pi$ for $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$.



(e) $hh\pi$ for $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$.

Figure 4.2: $hh\pi\pi$ and $hh\pi$ invariant mass distributions of $B^0 \rightarrow D^- \pi^+$, $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ and $B_s^0 \rightarrow D_s^- \rho^+$, with the $hh\pi\pi$ and $hh\pi$ mass windows overlaid as applicable.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

events are present in the samples for these to be reliably separated from $B_s^0 \rightarrow D_s^- \pi^+$ decays.

The $B^0 \rightarrow D^{*-} \pi^+$ background, which appears as peaking structures around the B^0 and D^{*-} masses in the $hh\pi\pi$ and $hh\pi$ invariant masses, respectively, is not excluded by the invariant mass windows, with the left-hand tail of the peak in $hh\pi$ extending below the upper limit of the $hh\pi$ window. However, unlike the signal, this background must proceed via a $D^0 \rightarrow h^- h^+$ resonance. A veto can therefore be constructed targeting the D^0 , in the form of exclusion windows of 30 MeV either side of the D^0 mass in the invariant masses of opposite-sign pairs of final-state hadrons, as was applied in Ref. [29]. Such windows were investigated for both the standard mass hypotheses and, in the cases of $K^+ K^-$ and $\pi^+ \pi^-$ pairs, with one of the particles assigned instead a π^\pm or K^\pm mass to target misidentified $D^0 \rightarrow K^- \pi^+$. The most effective windows, *i.e.*, those employed in the analysis, are $m(K^+ K^-) \notin [1830, 1890]$ MeV and $m(K^- \pi^+_{B_s^0}) \notin [1834, 1894]$ MeV for the $KK\pi\pi$ final state³, and windows of $m(\pi^- \pi^+_{B_s^0} [\text{as } K^+]) \notin [1834, 1894]$ MeV and $m(\pi^- \pi^+)_{\text{high}} \notin [1834, 1894]$ MeV for the $\pi\pi\pi\pi$ final state.

4.1.4 Studies of possible remaining backgrounds

Following the offline selection, the only backgrounds listed in Table 4.1 not directly addressed are the charmless and track-swap contributions. These contributions are not vetoed or modelled as they are assumed to be suppressed to a negligible level by the global offline selection requirements, though they must be studied further to validate such assumptions.

Firstly, the charmless background, is studied by examining the upper D_s^- sideband, *i.e.*, the region of the $hh\pi$ invariant mass distribution above the window applied in the selection, taken here to be $m(hh\pi) > 2010$ MeV. In this region, the only contributions to the $m(hh\pi\pi)$ which would be expected to peak would be $B^0 \rightarrow D^{*-} \pi^+$ (for both final states) and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ (for $KK\pi\pi$), as these both involve a three-body $hh\pi$ intermediate with a mass greater than that of the D_s^- , though neither peaks at the B_s^0 mass in $m(hh\pi\pi)$. A key difference between the charmless background component and signal is that, whilst the long lifetime of the D_s^- in signal ensures that the reconstructed vertex of the $hh\pi$ system is detached from the B_s^0 vertex per Section 2.2.4, no such separation of the reconstructed $hh\pi$ vertex from the B_s^0 vertex is present for the charmless background. This can be quantified using the flight-distance significance,

$$\text{FD}_{\text{sig.}}(D_s^-) = \frac{z_{\text{vtx.}}(D_s^-) - z_{\text{vtx.}}(B_s^0)}{\sqrt{\sigma_{z_{\text{vtx.}}}^2(D_s^-) + \sigma_{z_{\text{vtx.}}}^2(B_s^0)}}, \quad (4.7)$$

where $z_{\text{vtx.}}$ and $\sigma_{z_{\text{vtx.}}}$ are the central value and error, respectively, on the distance of a given vertex from the interaction point along the collision axis. Plotting $m(hh\pi\pi)$ with no $\text{FD}_{\text{sig.}}$ requirement, with $\text{FD}_{\text{sig.}} > 2$ and with $\text{FD}_{\text{sig.}} > 9$, as shown in Fig. 4.3, the only peaking structures observed were the two anticipated decays, both of which were removed when enforcing a $\text{FD}_{\text{sig.}}$ requirement. It was therefore concluded that no charmless background contributions passed the offline selection.

Considering now the contribution of track-swap background, $hh\pi$ invariant mass combinations were formed in which the associated pion is used in place of one of the hadrons, as shown in Fig. 4.4. In most cases, the entire $m(hh\pi)$ distribution lies far above the

³The $m(K^+ K^-)$ window is placed 4 MeV lower than the other windows as the left-hand tail of the D^0 peak is larger in $m(K^+ K^-)$.

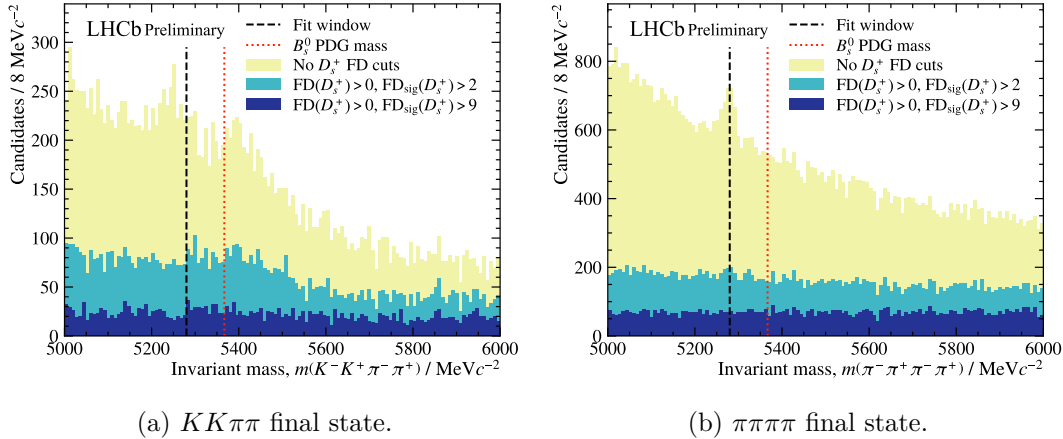
(a) $KK\pi\pi$ final state.(b) $\pi\pi\pi\pi$ final state.

Figure 4.3: $hh\pi\pi$ invariant masses of candidates with $m(hh\pi) > 2010$ MeV, subject to varying requirements on D_s^+ flight distance.

D^- and D_s^- masses; in the cases where candidates fell within this region, no peaking structures were observed, indicating that no candidates contained events wherein the associated pion originated from a D^- or D_s^- . It was thus concluded that the contribution of track-swap background is also negligible.

4.1.5 Replication of selection on MC samples

MC simulated samples of $B_s^0 \rightarrow D_s^- \pi^+$ and the remaining background contributions are used to model the components of the $hh\pi\pi$ invariant mass distribution, discussed in Section 4.3.1, and evaluate efficiencies and asymmetries, discussed in Section 4.2.1 & 4.4.1, respectively. To fulfil these purposes, the offline selection applied in Section 4.1.3 must be replicated on MC samples. Kinematic and topological variables, such as the $hh\pi\pi$ invariant mass, are sufficiently faithfully reproduced in MC simulation that selection applied to these variables in data may also be applied directly in the MC samples. This is the case for the invariant mass windows, $B^0 \rightarrow D^{*-} \pi^+$ veto and HCAL separation cut. The PID and L0 trigger responses, which have complex relations to the properties of each particle, are not necessarily faithfully reproduced in MC simulation, and thus cannot be applied directly. Instead, the per-event efficiencies of the PID and L0 requirements, described in Sections 4.2.2 & 4.2.3, respectively, are assigned as per-event weights in the MC samples.

4.2 Estimation of efficiencies

The efficiencies of kinematic selection, PID and L0 trigger requirements, and long track reconstruction must be understood in this analysis, as these are used for two purposes. Firstly, these efficiencies are required to construct the model of the $hh\pi\pi$ invariant mass used to determine A_{raw} (described in Section 4.3.1): the kinematic selection and PID efficiencies are required to determine the yields of background components from the $B_s^0 \rightarrow D_s^- \pi^+$ yield; the PID and L0 trigger efficiencies are incorporated as per-event weights in the MC samples to which template PDFs are fit (see Section 4.2.5). Secondly, all four efficiencies, split separately on the final state charge, are required to evaluate their respective contributions to the detection asymmetry in Section 4.4. Note that only the single-pion tracking asymmetry is determined directly from the corresponding efficiencies, hence only the pion tracking efficiency is discussed here. As each of the

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

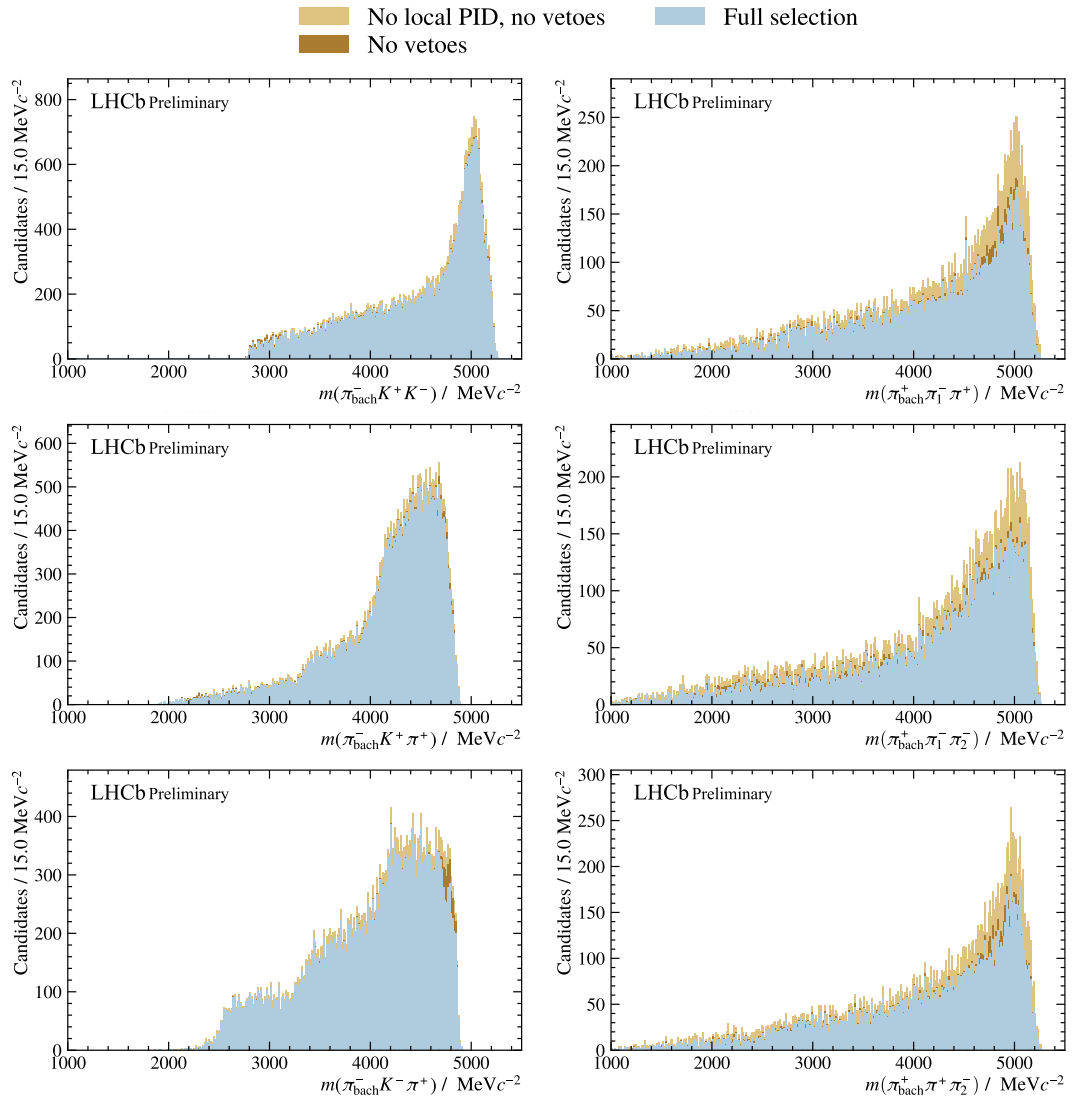


Figure 4.4: Invariant masses of the combination of the associated π^+ with each pair of hadrons from the D_s^- decay, for the $KK\pi\pi$ (left) and $\pi\pi\pi\pi$ (right) final states.

efficiencies arise from different sources, the efficiencies can be computed independent of one another, *e.g.*, the trigger efficiency may be reliably estimated in a sample in which the effect of PID requirements is not accounted for. The HLT1 efficiency is not included here as imposing the requirement of HLT1(Two)TrackMVA triggering on the B_s^0 removes only a handful of events, implying an efficiency of effectively 100%, given the offline selection and vetoes applied.

4.2.1 Kinematic selection efficiencies

The kinematic and topological properties of $B_s^0 \rightarrow D_s^- \pi^+$ decays are typically well-modelled in MC simulated samples. The efficiencies of kinematic selection requirements, *i.e.*, the invariant mass windows, HCAL separation cut, and $B^0 \rightarrow D^{*-} \pi^+$ veto, can therefore be calculated directly. Each efficiency is given by the fraction of MC candidates passing the selection:

$$\varepsilon_{\text{sel.}} = \frac{N_{\text{pass}}}{N_{\text{total}}}. \quad (4.8)$$

Whilst the HCAL separation requirement is imposed to enable the reliable estimation of L0 efficiencies, per Section 4.2.3, it is not accounted for in the L0 efficiencies themselves and thus must be accounted for in $\varepsilon_{\text{sel.}}$. The separation of each pair of final-state hadrons in each final state, that is, the distance between the projections of each of the final-state hadrons in the (x, y) plane of the HCAL, are shown in Fig. 4.5. A large fraction of the pairs lie below the threshold of 1000 mm; the like-sign pairs of hadrons originating from the D_s^- decays have the smallest separations. This is unsurprising as these hadrons are kinematically similar due to their common origin and are deflected in the same direction, hence these are likely to occupy nearby regions in the HCAL.

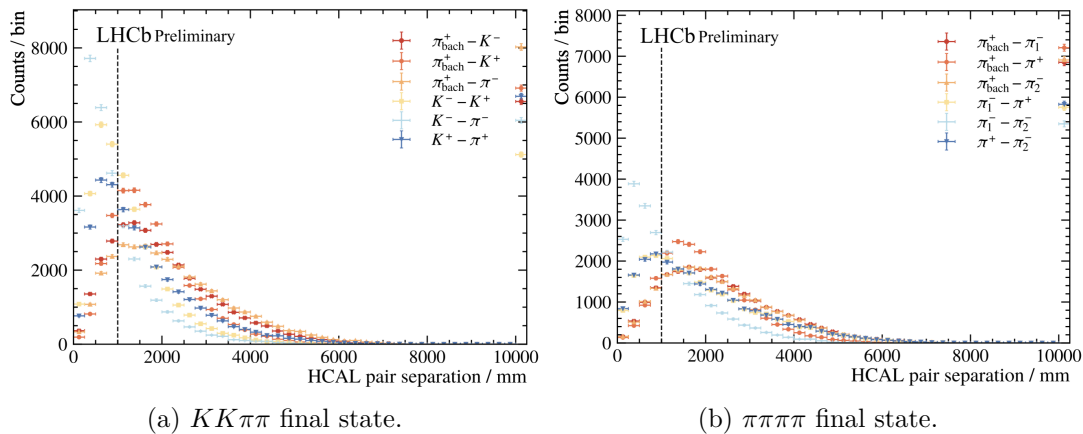


Figure 4.5: HCAL separation of each pair of final-state hadrons.

The efficiencies of the kinematic selections are computed in MC samples of $B_s^0 \rightarrow D_s^- \pi^+$ decays for $D_s^+ \rightarrow K^+ K^- \pi^+$ and $D_s^+ \rightarrow \pi^+ \pi^- \pi^+$, reflecting each year and magnet polarity of data-taking. The combination of these efficiencies weighted by the luminosity of their respective data-taking years and magnet polarities are listed in Table 4.3.

As anticipated, these efficiencies is largely dominated by the HCAL separation cut. In each final state, these efficiencies are in extremely good agreement, as is discussed further in Section 4.4.1.

Table 4.3: Efficiencies of kinematic selections, calculated with and without distinction on final-state charges.

Final state	Efficiency / %	
	$h^+ h^- \pi^+ \pi^-$	$h^- h^+ \pi^- \pi^+$
$KK\pi\pi$	24.067 ± 0.021	24.052 ± 0.021
$\pi\pi\pi\pi$	23.949 ± 0.081	23.937 ± 0.081

4.2.2 PID efficiencies

The quality of the `ProbNN` variables in MC simulated samples has a strong and complex dependence on the accuracy with which the interactions of final-state particles with the PID subdetectors, and their subsequent responses, are modelled. It is therefore extremely difficult to reliably extract efficiencies of PID requirements from MC in the same manner as those described in Section 4.2.1. Instead, a data-driven approach provides access to the PID efficiency without reliance on the accuracy of modelling in MC simulation. A standardised data-driven approach has been developed for estimating LHCb PID efficiencies in which reference tables are produced from dedicated calibration data samples in channels such as $D^{*+} \rightarrow D^0 \pi^+$, for a given set of PID requirements, which can then be applied to MC (reference) samples, implemented in the Python-based common software tool `PIDCalib2` [104]. The $D^{*+} \rightarrow D^0 \pi^+$ calibration samples are selected without PID requirements on the $D^0 \rightarrow K^- \pi^+$ decay products, with the π^+ of the D^{*+} used to tag the flavour of the D^0 , and thus the identities of the K^- and π^+ . The efficiency of a given PID requirement on the K^-/π^+ can then be evaluated per Eq. 4.8, specifically, in bins of the K^-/π^+ kinematics (typically p and η) with the calibration sample weighted such that it is kinematically similar to the reference sample. To obtain the number of signal events in the calibration sample before and after the PID requirement is applied, NLL fits to $m(K^- \pi)$ prior to applying the PID requirement are performed and the `sPlot` method [105] applied to obtain per-event `sWeights` which subtract contributions from combinatorial background. The sums of these events in each bin with and without the PID requirement provide the per-bin total and passing yields, respectively, from which the per-bin efficiencies are computed to produce a reference table.

To compute an efficiency for the whole candidate, the per-particle efficiencies of each final-state hadron are determined and the product of these taken:

$$\varepsilon_{\text{PID}}(h^+ h^- \pi^+ \pi^-) = \varepsilon_{\text{PID}}(h^+) \cdot \varepsilon_{\text{PID}}(h^-) \cdot \varepsilon_{\text{PID}}(\pi^+) \cdot \varepsilon_{\text{PID}}(\pi^-). \quad (4.9)$$

As two different PID requirements are applied for candidates in the resonant/non-resonant regions of the $D_s^+ \rightarrow h^+ h^- \pi^+$ phase-space (per Section 4.1.3), the $B_s^0 \rightarrow D_s^- \pi^+$ MC samples are split by phase-space region. In each phase-space region, the efficiency of the respective PID requirement is evaluated for the MC samples corresponding to each year and magnet polarity of data-taking, both per charge-conjugate final state and with no distinction on final state. Efficiencies per year/magnet polarity are combined by according to their luminosity; efficiencies per region are combined according to the yields of $B_s^0 \rightarrow D_s^- \pi^+$ obtained from invariant mass fits (per Section 4.3) to subsamples of the data samples in the respective $D_s^+ \rightarrow h^+ h^- \pi^+$ phase-space regions.

The charge-specific and charge-agnostic efficiencies for each phase-space region, combined over the data-taking years and magnet polarity, and their combination are listed in Table 4.4.

Table 4.4: Efficiencies of PID requirements, calculated with and without distinction on final-state charges.

Final state	Phase-space region	Efficiency / %	
		$h^+h^-\pi^+\pi^-$	$h^-h^+\pi^-\pi^+$
$KK\pi\pi$	Non-resonant	27.036 ± 0.082	27.056 ± 0.082
	Resonant	42.588 ± 0.058	42.564 ± 0.058
	Combination	40.050 ± 0.050	40.030 ± 0.050
$\pi\pi\pi\pi$	Non-resonant	34.33 ± 0.20	34.30 ± 0.20
	Resonant	53.99 ± 0.40	53.86 ± 0.40
	Combination	42.68 ± 0.21	42.61 ± 0.21

The efficiencies of the requirements on `ProbNNk` and `ProbNNpi` are relatively similar between the two final states. As expected, the efficiencies outside the resonant regions are notably lower; however, this is deemed to be necessary to provide an improved signal purity for non-resonant candidates. The greater fraction of resonant candidates in the $KK\pi\pi$ final state (shown in Fig. 4.1) is reflected in a combined efficiency much closer to the efficiency for resonant candidates than in the $\pi\pi\pi\pi$ final state. As for the kinematic selection efficiencies, the PID efficiencies are in very good agreement between charge final states, as discussed later in Section 4.4.2.

4.2.3 L0 trigger efficiencies

Similarly to the PID requirements, the response of the L0 trigger, or more specifically `L0Hadron`, cannot be reliably obtained from MC, as it depends on the modelling of the interactions of hadrons with the HCAL, ECAL, SPD and PS. The `L0Hadron` efficiency must also be estimated by a data-driven method. The $D^{*+} \rightarrow D^0\pi^+$ calibration samples recorded for PID efficiency determination are used to access the efficiency of `L0Hadron`. Examining the K^- and π^+ of the $D^0 \rightarrow K^-\pi^+$ decay individually, requirements are imposed that each hadron has fewer than 450 corresponding hits in the SPD and that the `L0Muon` trigger line fires independently of (TIS, discussed in Section 5.2.2) the hadron. This ensures that these samples are statistically independent of whether `L0Hadron` has triggered on said hadron. The `L0Hadron` efficiency is then constructed as the fraction of these candidates for which the `L0Hadron` triggers on the hadron of interest. As for the PID efficiencies, per-event *sWeights* are obtained from fits to $m(K^-\pi)$, and used to obtain yields in bins of hadron E_T (rather than p_T , as this is the quantity measured by the HCAL) before and after requiring `L0Hadron` to trigger on the hadron computed by fitting to $m(K^-\pi)$. These efficiencies are computed separately for the decays of D^{*+} and D^{*-} , and for hadrons in the inner and outer regions of the HCAL.

This approach works well if the final-state hadrons are well separated in the HCAL; however, if two (or more) final-state hadrons deposit energy sufficiently close together, the energy deposits are merged and the `L0Hadron` efficiency is subsequently higher. To account for this, efficiency tables are computed solely from $K^-\pi^+$ pairs with a separation distance, d , *i.e.*, the distance between their projections onto the HCAL (x, y) plane, of at least 1000 mm, with both final-state hadrons within the HCAL acceptance. At this separation, the effect of merged energy deposits is negligible.

These efficiencies are extrapolated to poorly-separated hadron pairs and pairs with one or both hadron outside the HCAL acceptance by a parameterisation of the efficiency.

This parameterisation is defined as a function of the fraction of tracks which do not extend to the HCAL, the E_T of the K^- and π^+ , and the calibration of the HCAL in each cell. This provides efficiencies for pairs partially within the HCAL acceptance approximately consistent with the least well-separated pairs passing the separation cut, *i.e.*, those separated by ~ 1000 mm. However, for poorly separated pairs of hadrons ($d < 1000$ mm), the `L0Hadron` efficiency increases rapidly. As this rapid increase would enhance the effect of any inaccuracies on the `L0Hadron` efficiencies, and by extension the `L0Hadron` contribution to $A_{\text{det.}}$ discussed in Section 4.4.3, candidates containing any poorly separated pair of hadrons were rejected. This requirement has a low signal efficiency (see Table 4.3), respectively, per Table 4.3, but is essential if the efficiencies of and asymmetries due to `L0Hadron` are to be accurately estimated.

As candidates are accepted if `L0Hadron` triggers on any of the final-state hadrons, *i.e.*, candidates are rejected only if `L0Hadron` triggers on none of the final-state hadrons, the per-event efficiency can be computed from per-particle efficiencies as

$$\varepsilon_{\text{L0,TOS}} = 1 - \prod_i (1 - \varepsilon_{\text{L0,TOS}}^i). \quad (4.10)$$

This differs from the equivalent for the efficiency of the PID requirements, in Eq. 4.9, as all, rather than any, of the hadrons must pass the PID requirements.

The full approach described above is implemented in a Python-based common software tool, `L0HadronTables` [106, 107], which applies the efficiency reference tables per year/magnet polarity of data-taking to MC simulated signal samples. This tool is applied to the $B_s^0 \rightarrow D_s^- \pi^+$ MC samples reflecting each year and magnet polarity of data-taking, with the efficiencies (combining across year and magnet polarity according to respective luminosities) listed in Table 4.5. It is particularly important that these are computed per data-taking year, as the E_T thresholds in the L0 trigger were modified each year, as reported in Ref. [62].

Table 4.5: Efficiencies of `L0Hadron`, calculated with and without distinction on final-state charges.

Final state	Efficiency / %	
	$h^+ h^- \pi^+ \pi^-$	$h^- h^+ \pi^- \pi^+$
$KK\pi\pi$	39.0610 ± 0.031	39.193 ± 0.031
$\pi\pi\pi\pi$	41.54 ± 0.12	41.54 ± 0.12

Similar to both the kinematic selection and PID efficiencies, the `L0Hadron` efficiencies are of a comparable scale between final states. Notably, the charge-specific efficiencies evaluated in $K^+ K^- \pi^+ \pi^-$ and $K^- K^+ \pi^- \pi^+$ are not in agreement, as discussed later in Section 4.4.3.

4.2.4 Single-pion tracking efficiencies

The efficiency to reconstruct the trajectories of charged pions as long tracks is required to evaluate the contribution of differences in these efficiencies for π^+ and π^- to $A_{\text{det.}}$, as described in Section 4.4.4. This can be (and traditionally was) evaluated, to a reasonable accuracy, from MC simulated samples akin to the method applied in Section 4.2.1. However, the systematic uncertainties assigned to these efficiencies from differences between data and MC proved to be large in early LHCb measurements [108]. A data-driven method was therefore developed to circumvent these effects, in which muons

from $J/\psi \rightarrow \mu^+\mu^-$ decays reconstructed as tracks of given types (per Fig. 3.5 in Section 3.2.1) are compared to estimate the efficiency to reconstruct VELO, T and long tracks, as described in Ref. [109]. This method assumes that the track reconstruction of hadrons differs from that of muons only by their interactions with the detector material, and that these effects factor out from these efficiencies. As these interactions are assumed to be charge-symmetric (to the scale of $\mathcal{O}(10^{-4})$ for pions [31]), and as only the charge-dependent effects in the long track reconstruction efficiencies are of interest here, these effects are neglected and assumed to cancel in the contributions to A_{det} in the $\pi\pi\pi\pi$ final state.

To determine the efficiency to reconstruct a muon as a given track type, dedicated samples of $J/\psi \rightarrow \mu^+\mu^-$ decays are collected, where the J/ψ is required to arise from a b -hadron decay to improve the purity of the samples. One of the μ^\pm , dubbed the “tag” muon, is reconstructed as a long track, whilst the other, dubbed the “probe” muon/track is reconstructed using information from all but at least one of the tracking systems, *i.e.*, as a VELO, upstream, downstream or T track. The probe tracks are compared to fully reconstructed tracks (those containing information from the subdetectors of the probe track and the missing subdetector(s)), with the efficiency of the subdetector(s) in the track reconstruction then given by the fraction of probe tracks for which a corresponding fully reconstructed track is found. As long tracks require hits in the VELO and T1-3 but not the TT and muon systems, probe tracks for studying the long track reconstruction efficiency are formed from hits in the TT and the muon systems using a dedicated reconstruction. Of the possible reconstructed long tracks in the event, if the extrapolation of any long track into the muon systems is compatible with at least 70% of the probe track hits in the muon station (and at least 60% of hits in the TT if applicable), then the probe track is considered to correspond to a long track. The number of probe tracks before and after matching to the long tracks are obtained, in bins of the track p_T and η , by NLL fits to $m(\mu^+\mu^-)$, from which the efficiency is then computed. The application of this method in this analysis is discussed later, in Section 4.4.4.

4.2.5 Efficiencies as per-event weights

As discussed in Section 4.1.5, the PID and L0 requirements cannot be reliably applied directly to the MC samples. Instead, the per-event efficiencies of each of these requirements are estimated per Sections 4.2.2 & 4.2.3, with the combined efficiency taken as the product of the two. Specifically, the PID efficiency is applied for all samples, whilst the L0 requirement, which relies on dedicated information included which is only included in the $B_s^0 \rightarrow D_s^- \pi^+$ sample, is only incorporated for $B_s^0 \rightarrow D_s^- \pi^+$. Since the $hh\pi\pi$ and $hh\pi$ invariant mass distributions are not significantly modified by the L0 requirement, the impact of not including the corresponding efficiencies is assumed to be negligible.

Applying these as per-event weights, the $KK\pi\pi$ and $KK\pi$ invariant masses, and the p_T and η of the associated π^+ are shown in Fig. 4.6. Further distributions are shown in Figs. A.1 & A.2 in Appendix A.2.

Whilst the invariant mass distributions undergo a small change in resolution, becoming marginally narrower, the largest modifications from the inclusion of these efficiencies are in the final-state hadron kinematics. Understanding the effect of these weights on the kinematics of the final state hadrons are crucial to accurately determining their corresponding detection asymmetries.

The L0 weights introduce a double-peak structure in the p_T distribution of each final-state hadron. The peak at higher p_T is driven by the efficiency to trigger on the associated

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

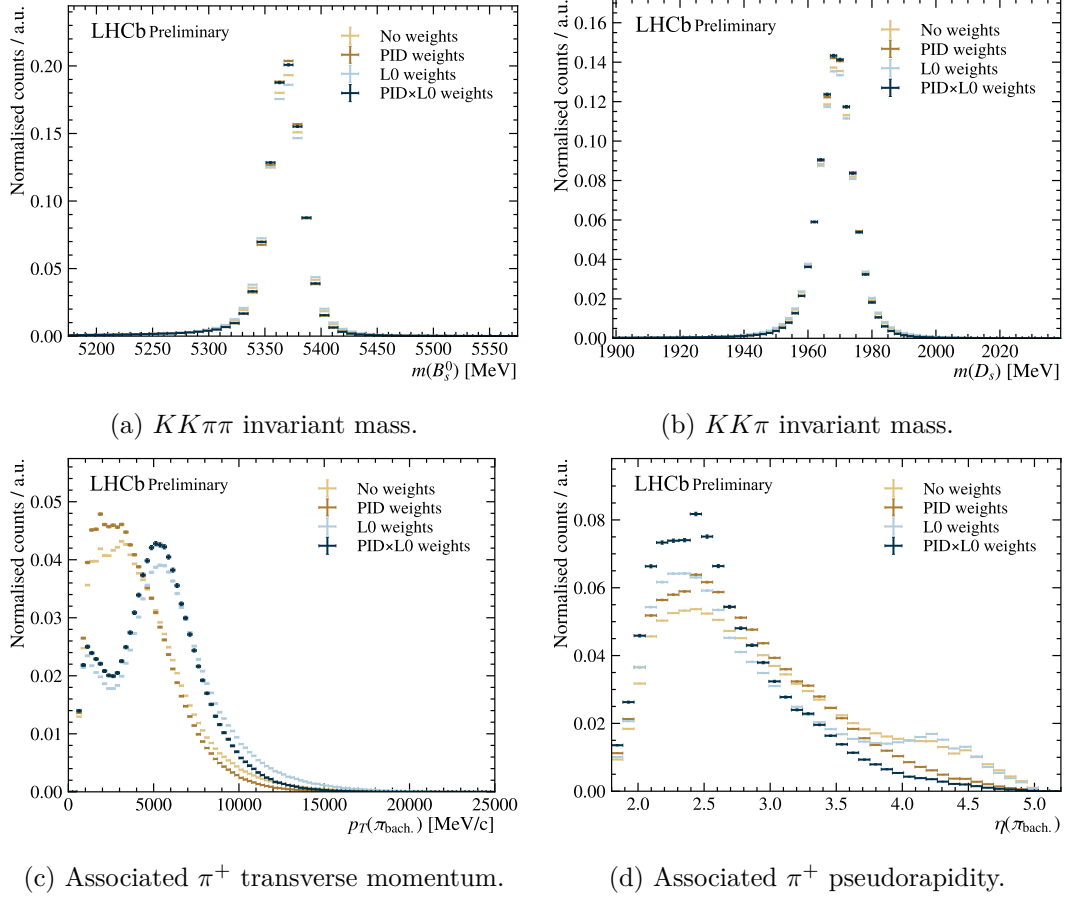


Figure 4.6: Effects of per-event weights from `PIDCa1ib2` and `L0HadronTables`, and their combination, on key distributions in the $KK\pi\pi$ final state.

π^+ , which increases with increasing p_T . The peak at lower p_T is driven by the efficiency to trigger on the other hadron(s) in the candidate; since the kinematics of the final-state hadrons are constrained by the B_s^0 and D_s^- , events in which the other hadron(s) have a greater p_T (and consequently a higher efficiency) will generally contain associated π^+ with a lower p_T .

The PID efficiencies generally have a smaller effect; however, they are responsible for a significant reduction in candidates with final-state hadron at high η .

4.3 Raw asymmetry, A_{raw}

The raw asymmetry, given in Eq. 4.6, is computed from the signal yields in each charge-conjugate final state. These charge-specific signal yields, parameterised by A_{raw} and a total signal yield, are obtained by a 1-dimensional NLL fit [110] to the $hh\pi\pi$ invariant mass, performed simultaneously over the charge-conjugate final states $h^-h^+\pi^-\pi^+$ and $h^+h^-\pi^+\pi^-$. Each of component forming the $hh\pi\pi$ invariant mass spectrum is modelled by a probability density function (PDF), with the shape taken from fits to respective MC simulated samples. The PDFs chosen were informed largely by Ref. [29], which performs a similar fit, albeit in 2-dimensions: over the B_s^0 mass (computed with the constraint of the D_s^- mass taking its measured value per Ref. [5] using `DecayTreeFitter (DTF)` [111]) and $hh\pi$ invariant mass, the latter of which is not necessary in this analysis as, per

Section 4.1.4, the contribution of charmless background is negligible in the selected samples.

4.3.1 Modelling of signal and background components

The PDF chosen to model the reconstructed $B_{(s)}^0$ mass in $B_{(s)}^0 \rightarrow D_{(s)}^- h^+$ decays, must account for two effects. Firstly, they must account for the detector resolution, which results in a Gaussian-like distribution of invariant masses. Secondly, they must account for energy losses from final-state radiation, which result in an extended left-hand tail, often modelled by a power-law distribution. Three PDFs are considered to model the $hh\pi\pi$ invariant mass of $B_s^0 \rightarrow D_s^- \pi^+$ decays: the double-sided Crystal Ball (DSCB), the sum of an extended Cruijff and a Gaussian, and the sum of a double-sided Hypatia⁴ and Johnson S_U function, which are defined according to Eqs. A.1, A.4 and A.3/A.7, respectively. Each of these were fit to the $B_s^0 \rightarrow D_s^- \pi^+$ MC samples in each final state (with ProbNN and L0 efficiencies applied as per-event weights), as shown in Fig. 4.7.

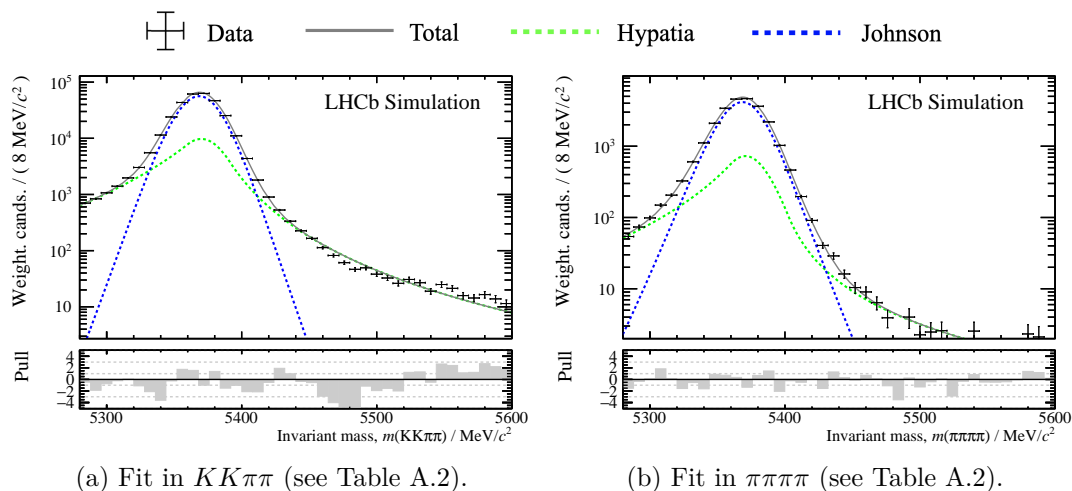


Figure 4.7: Fits of the Hypatia+Johnson S_U model to the 2016-18 combined $B_s^0 \rightarrow D_s^- \pi^+$ MC samples in both $hh\pi\pi$ final states. Equivalent fits using the DSCB and extended Cruijff+Gaussian models are shown in Fig. A.3 in Appendix A. The corresponding parameters are listed in Table A.2.

Both the Extended Cruijff+Gaussian and Hypatia+Johnson S_U PDFs provide a better description of $m(hh\pi\pi)$ for $B_s^0 \rightarrow D_s^- \pi^+$ than the DSCB. The Extended Cruijff provides a better description of the right-hand tail, whilst the Hypatia+Johnson S_U PDF provides a better description of the core of the distribution. As the majority of events lie within the core, the Hypatia+Johnson S_U model was chosen to describe the $B_s^0 \rightarrow D_s^- \pi^+$ component.

Considering each background component, *i.e.*, those in Table 4.1 not marked by * or †, many occupy similar regions of $m(hh\pi\pi)$. Whilst this makes such backgrounds highly-correlated in the fit and hence difficult to separate, if the backgrounds are sufficiently similar then they may be safely described by a single component. In particular, the $B^0 \rightarrow D_s^+ \pi^-$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$ backgrounds both lie at $m(hh\pi\pi) < 5300 \text{ MeV}/c^2$, and the $B_s^0 \rightarrow D_s^- K^+$ and $B^0 \rightarrow D^- \pi^+$ backgrounds both lie beneath the signal peak, centred

⁴It is henceforth implied that Hypatia refers to the double-sided Hypatia function.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

around $m(hh\pi\pi) \sim 5350 \text{ MeV}/c^2$, as shown in Fig. 4.8. The former pair are sufficiently similar that these can be modelled by a single PDF; the component representing both $B^0 \rightarrow D_s^+ \pi^-$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$ is referred to throughout as the “low mass” background, and is described by the $B^0 \rightarrow D_s^+ \pi^-$ shape, which can be taken as the $B_s^0 \rightarrow D_s^- \pi^+$ model with the B^0 - B_s^0 mass difference subtracted from the mean parameter, μ . The latter pair are insufficiently similar to model together and require a careful treatment of their yields/asymmetries to mitigate correlations both between the two backgrounds and between these backgrounds and $B_s^0 \rightarrow D_s^- \pi^+$ signal.

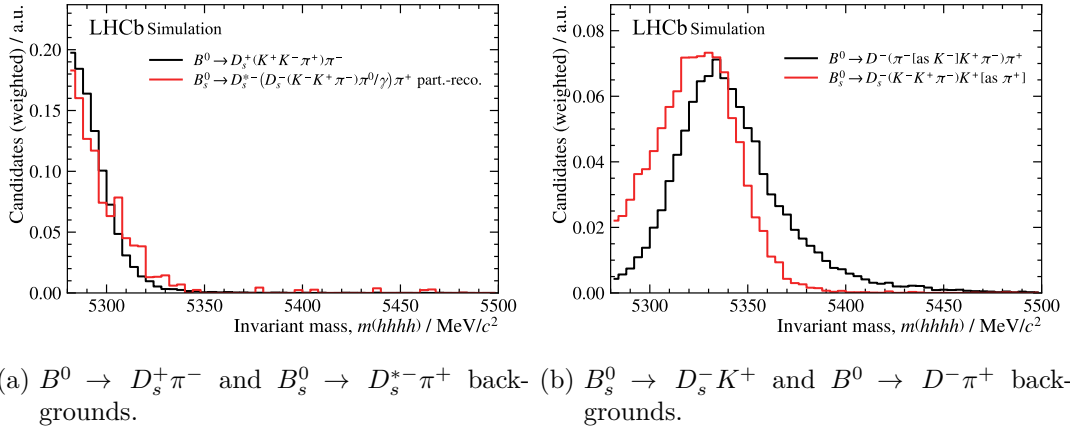


Figure 4.8: Comparison of peaking background components occupying similar spaces in the $KK\pi\pi$ invariant mass.

For backgrounds with misidentified final-state particles ($B_s^0 \rightarrow D_s^- K^+$, $B^0 \rightarrow D^- \pi^+$ and $A_b^0 \rightarrow \bar{A}_c^- \pi^+$), these can be modelled empirically by finding a PDF which describes these suitably in MC, modified to the $hh\pi\pi$ mass hypothesis. The Hypatia+Johnson S_U model was trialled as a model for these: in $B_s^0 \rightarrow D_s^- K^+$, the full Hypatia+Johnson S_U model was required for an adequate description, whilst for $B^0 \rightarrow D^- \pi^+$ and $A_b^0 \rightarrow \bar{A}_c^- \pi^+$, the Johnson S_U alone (described in Eq. A.7) was sufficient. Fits of these models to their respective MC samples are shown in Fig. 4.9. As the $B_s^0 \rightarrow D_s^- K^+$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$ MC samples were only available in the $KK\pi\pi$ final state, and since their shape in $m(hh\pi\pi)$ is independent of the decay mode of the D_s^- , the $KK\pi\pi$ models of these backgrounds were also used in the $\pi\pi\pi\pi$ final state.

Finally, the combinatorial background in hadronic b -meson decays is typically modelled by exponential-like terms, *e.g.*, the double-exponential model used in Ref. [29]. In this measurement, the offline selection described in Section 4.1 leaves comparatively little combinatorial background, particularly in the $KK\pi\pi$ final state. This leaves insufficient information to reliably fit the 3 degrees of freedom of the double-exponential model (two exponents and a relative scale factor). Instead, an exponential+constant model, *i.e.*, a double-exponential model with one exponent fixed to 0, is compared to a single exponential model. To provide sufficient combinatorial background to reliably fit these models, dedicated samples of data were prepared in which the stricter ProbNN requirements in the non-resonant $D_s^+ \rightarrow h^+ h^- \pi^+$ phase space and the $m(hh\pi)$ window are not applied. The combinatorial background of these dedicated sample and of the fully selected samples was compared and deemed to be sufficiently similar to be used to obtain the parameters of the combinatorial models.

The two models were fit to the dedicated samples, specifically to the upper $hh\pi\pi$ side-

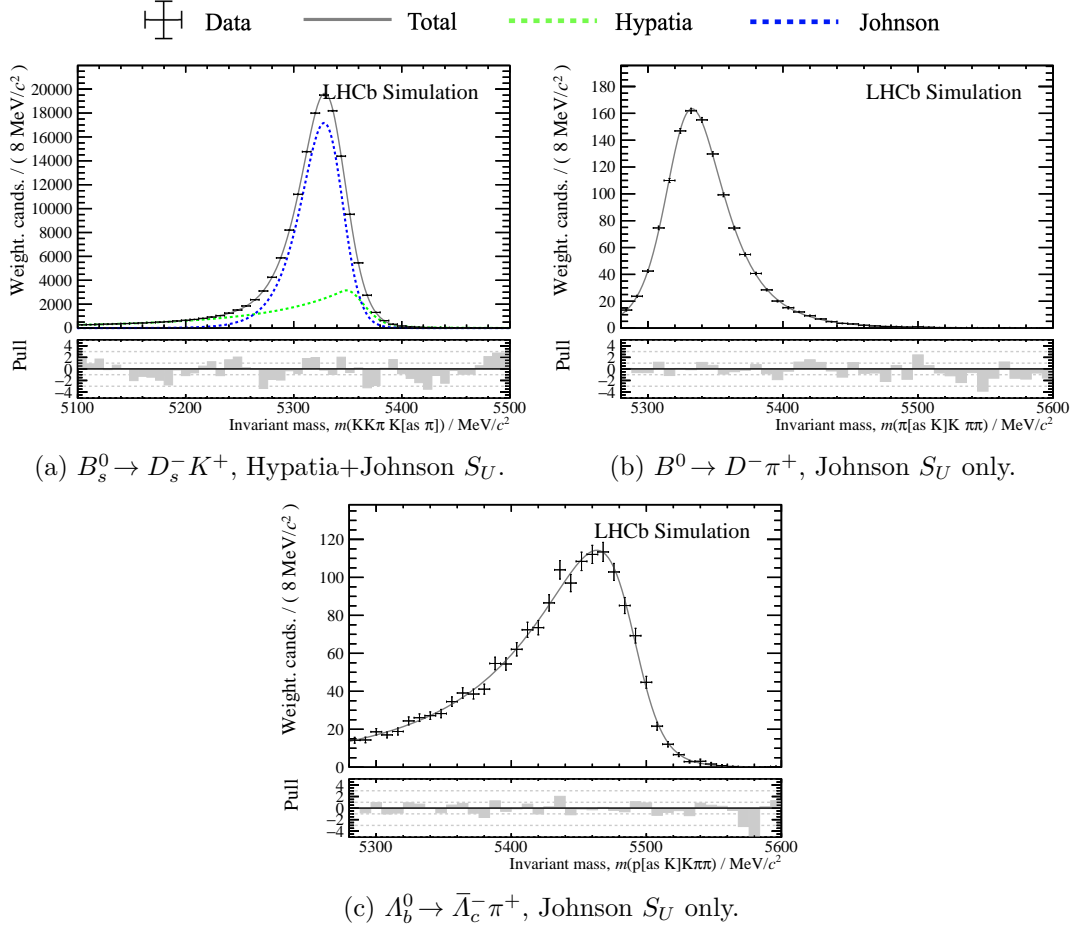


Figure 4.9: Template PDFs modelling each peaking background contribution.

band, $m(hh\pi\pi) \in [5600, 6800] \text{ MeV}/c^2$, as shown in Fig. 4.10. Both models describe the sideband well; however, the exponential+constant model was chosen as the constant term provides a sizeable fraction of the combinatorial background shape in the flatter region of high $m(hh\pi\pi)$. This allows the exponential term to account for a steeper combinatorial background shape for lower values of $m(hh\pi\pi)$, closer to $m_{B_s^0}$, in the fits to data.

With all components modelled, the total $hh\pi\pi$ invariant mass model is constructed as a sum of the PDFs, each extended by a corresponding yield for each charge-conjugate final state, *i.e.*,

$$f_{\text{tot.}}(m) = \sum_i N_i^+ f_i(m|h^- h^+ \pi^- \pi^+) + N_i^- f_i(m|h^+ h^- \pi^+ \pi^-), \quad (4.11)$$

for each component, i . The fit, referred to throughout as the nominal fit, is therefore performed simultaneously over the two charge-conjugate final states, with the parameters of f_i shared between the states and the yields N_i^\pm separated. The charge-specific yields, N_i^\pm are parameterised in terms of a total, charge-agnostic yield N_i and an asymmetry A_i , as

$$N_i^\pm = \frac{1 \mp A_i}{2} N_i, \quad (4.12)$$

such that A_{raw} (which is equivalent to $A_{B_s^0 \rightarrow D_s^- \pi^+}$), can be extracted directly from the fits to data. This parameterisation also results in a more stable fit, allowing for direct

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

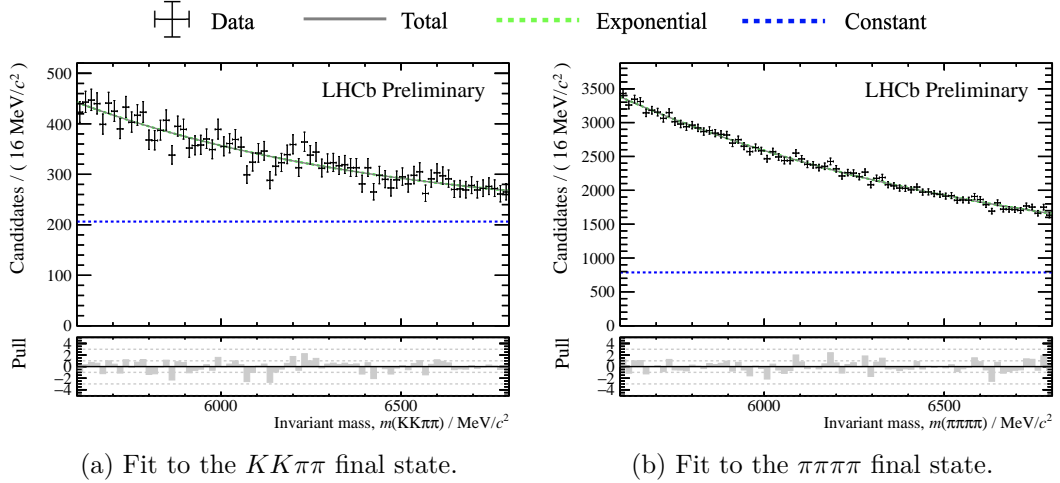


Figure 4.10: Fits of the exponential+constant model to the combinatorial background data samples (2016-18, both magnet polarities). The equivalent fits of a single exponential model are shown in Fig. A.4 in Appendix A.

control of the asymmetry of components in the subsequent pseudoexperiments performed to validate the fit and estimate systematic uncertainties. Furthermore, this enables a direct comparison to an equivalent charge-agnostic fit.

The parameters of each f_i are fixed in the nominal fit; however, the position and resolution of both signal and peaking background distributions in $m(hh\pi\pi)$ may not be modelled accurately by the MC. The fixed mean-like (μ in Eqs. A.3 & A.7), μ_{MC}^i , and width-like (σ in Eq. A.3 and λ in Eq. A.7), σ_{MC}^i , parameters of each model are thus modified by floating parameters in fits to data,

$$\mu_{\text{data}}^i = \mu_{\text{MC}}^i + \delta_\mu, \quad (4.13a)$$

$$\sigma_{\text{data}}^i = \sigma_{\text{MC}}^i + \delta_\sigma, \quad (4.13b)$$

where $\delta_\mu \in [-2.5, 2.5] \text{ MeV}/c^2$ and $\delta_\sigma \in [-5, 10] \text{ MeV}/c^2$. Similarly, the exponent of the exponential term in the combinatorial model is allowed to float in the fit to data, starting from its value in the fit to the dedicated sample. Meanwhile, the fraction between the exponential and constant terms of the combinatorial model is fixed; a corresponding systematic uncertainty is determined in Section 4.6.1 by comparing the resulting value of A_{fs}^s from the nominal fit and equivalent fit wherein this fraction is allowed to float under a Gaussian constraint.

The asymmetries of the largest components, $B_s^0 \rightarrow D_s^- \pi^+$, and the combinatorial and low mass backgrounds, are allowed to float, whilst the asymmetries of $B_s^0 \rightarrow D_s^- K^+$, $B^0 \rightarrow D^- \pi^+$ and $A_b^0 \rightarrow \bar{A}_c^- \pi^+$ are fixed to 0, as these are small components so any asymmetry here (not anticipated to be greater than $\mathcal{O}(1\%)$) induces a very small difference in the number of events in each charge-conjugate final state. This approach is validated by pseudoexperiments with asymmetries assigned in these backgrounds, discussed in Section 4.3.4. Systematic uncertainties are determined similarly to that of the combinatorial fraction, as discussed in Section 4.6.1.

The yields of $B_s^0 \rightarrow D_s^- \pi^+$, $A_b^0 \rightarrow \bar{A}_c^- \pi^+$, and the combinatorial and low mass backgrounds are allowed to float; however, as the $B_s^0 \rightarrow D_s^- K^+$ and $B^0 \rightarrow D^- \pi^+$ yields are smaller and their components overlap significantly with $B_s^0 \rightarrow D_s^- \pi^+$, their values relative to the

$B_s^0 \rightarrow D_s^- \pi^+$ yield must be fixed in the fit to data. These yields are parameterised in terms of $N_{B_s^0 \rightarrow D_s^- \pi^+}$ and a relative fraction, $f_{B_s^0 \rightarrow D_s^- \pi^+}^i$, as

$$N_i = f_{B_s^0 \rightarrow D_s^- \pi^+}^i \cdot N_{B_s^0 \rightarrow D_s^- \pi^+}. \quad (4.14)$$

This relative fraction, which is fixed in the nominal fit, is estimated by comparing the production rate (hadronisation and branching fractions) and efficiency of a given background with those of $B_s^0 \rightarrow D_s^- \pi^+$. It is assumed that the only components of the efficiency which differ between $B_s^0 \rightarrow D_s^- \pi^+$ and the backgrounds are those of the kinematic selections (mass windows and $B^0 \rightarrow D^{*-} \pi^+$ veto), and particle identification requirements, described in Sections 4.2.1 & 4.2.2. The tracking, trigger and other offline selection efficiencies are assumed to be sufficiently similar between $B_s^0 \rightarrow D_s^- \pi^+$ and each background that these need not be accounted for. Therefore, the relative fraction can be written as

$$f_{B_s^0 \rightarrow D_s^- \pi^+}^i = \frac{f_d}{f_s} \cdot \underbrace{\frac{\mathcal{B}(D^+ \rightarrow K^- \pi^+ \pi^+)}{\mathcal{B}(D_s^+ \rightarrow K^+ K^- \pi^+)}}_{\text{for } i=B^0 \rightarrow D^- \pi^+ \text{ only}} \cdot \frac{\mathcal{B}(i)}{\mathcal{B}(B_s^0 \rightarrow D_s^- \pi^+)} \cdot \frac{\varepsilon_{\text{kin.}}^i}{\varepsilon_{\text{kin.}}^{B_s^0 \rightarrow D_s^- \pi^+}} \cdot \frac{\varepsilon_{\text{PID}}^i}{\varepsilon_{\text{PID}}^{B_s^0 \rightarrow D_s^- \pi^+}}, \quad (4.15)$$

where $\frac{f_d}{f_s}$ (or $1/\frac{f_s}{f_d}$) is the ratio of the fragmentation fractions of B^0 and B_s^0 , $\mathcal{B}(i)$ are the branching fractions of process i , and $\varepsilon_{\text{kin.}}$ and ε_{PID} are as described in Section 4.2. The values of $f_{B_s^0 \rightarrow D_s^- K^+}^{B_s^0 \rightarrow D_s^- K^+}$ and $f_{B_s^0 \rightarrow D_s^- \pi^+}^{B^0 \rightarrow D^- \pi^+}$ in the 2018 magnet up configuration, and its constituent values, are stated in Table 4.6. As for the combinatorial exponent and fixed asymmetries, a systematic uncertainty corresponding to these fractions is determined by comparing against a fit in which these are floating these within a Gaussian constraint.

Table 4.6: Fractions describing the yields of $B_s^0 \rightarrow D_s^- K^+$ and $B^0 \rightarrow D^- \pi^+$ relative to the yield of $B_s^0 \rightarrow D_s^- \pi^+$ (per Eq. 4.15) and its constituent components. Kinematic and PID efficiencies are taken from MC corresponding to 2018 magnet up data-taking.

Quantity	Decay	Value	Source
Kinematic efficiency	$B_s^0 \rightarrow D_s^- \pi^+$	0.94155 ± 0.00085	—
	$B_s^0 \rightarrow D_s^- K^+$	0.70423 ± 0.00085	—
	$B^0 \rightarrow D^- \pi^+$	0.46239 ± 0.00074	—
PID efficiency	$B_s^0 \rightarrow D_s^- \pi^+$	$(3.821356 \pm 0.000001) \times 10^{-1}$	—
	$B_s^0 \rightarrow D_s^- K^+$	$(2.507092 \pm 0.000024) \times 10^{-2}$	—
	$B^0 \rightarrow D^- \pi^+$	$(2.72802 \pm 0.00013) \times 10^{-3}$	—
Branching fraction	$B_s^0 \rightarrow D_s^- \pi^+$	$(2.98 \pm 0.14) \times 10^{-3}$	[5]
	$B_s^0 \rightarrow D_s^- K^+$	$(2.25 \pm 0.12) \times 10^{-4}$	[5]
	$B^0 \rightarrow D^- \pi^+$	$(2.51 \pm 0.08) \times 10^{-3}$	[5]
	$D_s^+ \rightarrow K^+ K^- \pi^+$	$(5.37 \pm 0.01) \times 10^{-2}$	[5]
	$D^+ \rightarrow K^- \pi^+ \pi^+$	$(9.38 \pm 0.16) \times 10^{-2}$	[5]
Had. fraction, f_s/f_d	—	0.2539 ± 0.0079	[112]
Yield fraction	$B_s^0 \rightarrow D_s^- K^+$	0.00423 ± 0.00030	—
	$B^0 \rightarrow D^- \pi^+$	0.0203 ± 0.0014	—

4.3.2 The nominal invariant mass fit

The complete model described in Eq. 4.11 was fit to the full data samples (2016-18, magnet up and down combined) in the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states. The resulting distributions of these fits are shown in Fig. 4.11, with their corresponding parameters listed in Table 4.7 and the correlations between these parameters shown in Fig. 4.12.

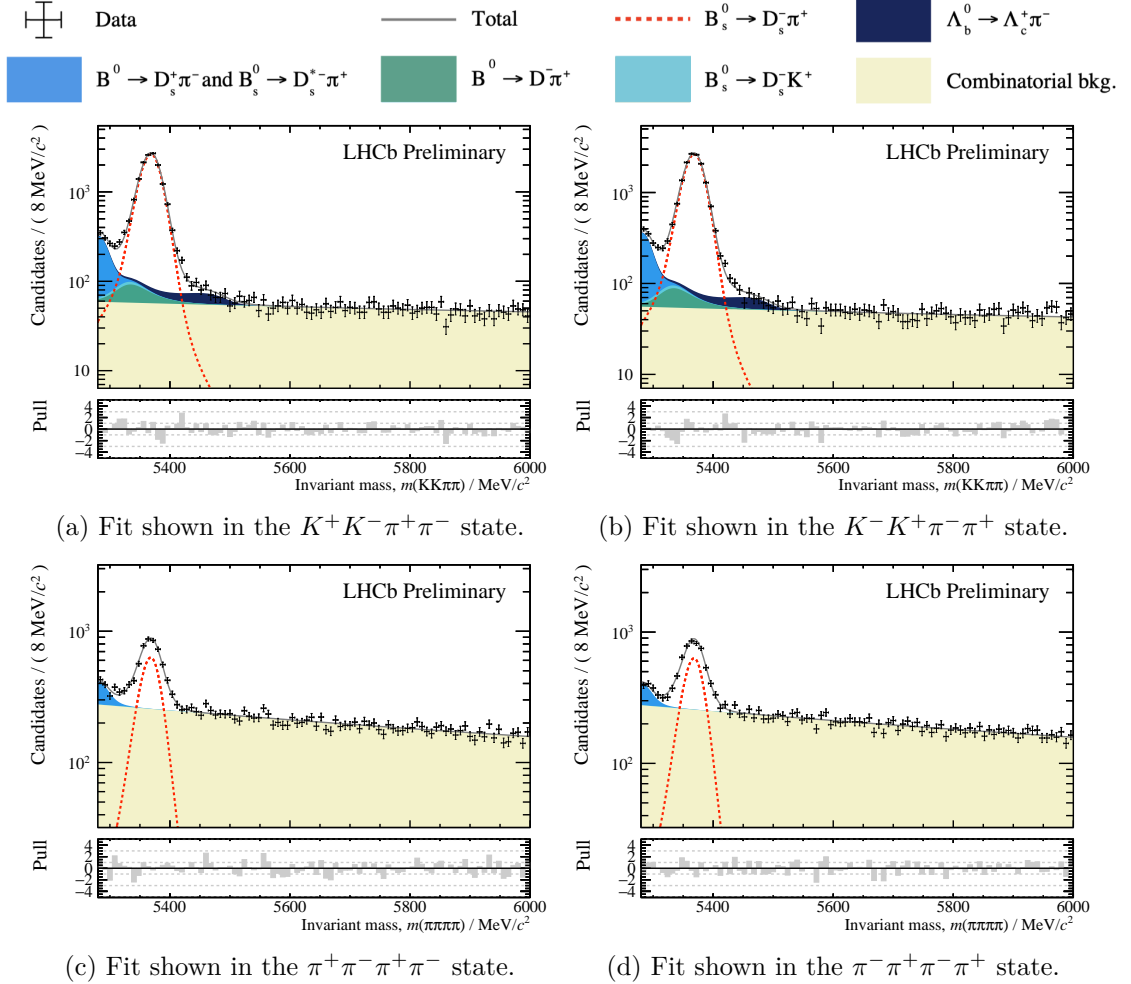


Figure 4.11: Nominal invariant mass fit to each final state, shown for each charge conjugate state. Corresponding parameters are listed in Table 4.7. The same fits are shown on a linear scale in Fig. A.5 in Appendix A.

The fits converge well, accurately describing the data in each final state and each charge conjugate state. Both fits return values of A_{raw} consistent with 0.

In the $KK\pi\pi$ final state, the combinatorial and low mass background asymmetries are both non-zero, to significances of 2.5σ and 2.4σ , respectively. In the combinatorial background, this appears to be physical and is visible in the differences in yields at high $m(KK\pi\pi)$. The low mass background asymmetry is more difficult to discern between a physical effect or an artefact of the model construction, as the component is anti-correlated with the combinatorial background, so may be driven by the combinatorial asymmetry, and the component is highly dependent on the statistical fluctuations at low $m(KK\pi\pi)$. In the $\pi\pi\pi\pi$ final state, the lower statistics and consequent lower sensitivity

Table 4.7: Parameters of the nominal invariant mass fits shown in Fig. 4.11.

Parameter	Fit value	
	$KK\pi\pi$ final state	$\pi\pi\pi\pi$ final state
$B_s^0 \rightarrow D_s^- \pi^+$ yield	28940 ± 190	7270 ± 150
Combinatorial yield	8960 ± 180	37620 ± 250
Low mass yield	1510 ± 57	720 ± 68
$A_b^0 \rightarrow \bar{A}_c^- \pi^+$ yield	600 ± 110	—
$B_s^0 \rightarrow D_s^- \pi^+$ asymmetry	$(-2.5 \pm 6.2) \times 10^{-3}$	$(-0.3 \pm 1.6) \times 10^{-2}$
Combinatorial asymmetry	$(-3.0 \pm 1.2) \times 10^{-2}$	$(2.8 \pm 5.4) \times 10^{-3}$
Low mass asymmetry	$(8.5 \pm 3.5) \times 10^{-2}$	$(-5.4 \pm 8.1) \times 10^{-2}$
Comb. exponent / $(\text{MeV}/c^2)^{-1}$	$(-9.0 \pm 2.6) \times 10^{-4}$	$(-1.187 \pm 0.057) \times 10^{-3}$
Mass value shift / MeV/c^2	-0.63 ± 0.12	-0.93 ± 0.36
Mass resolution shift / MeV/c^2	8.30 ± 0.39	3.41 ± 0.82

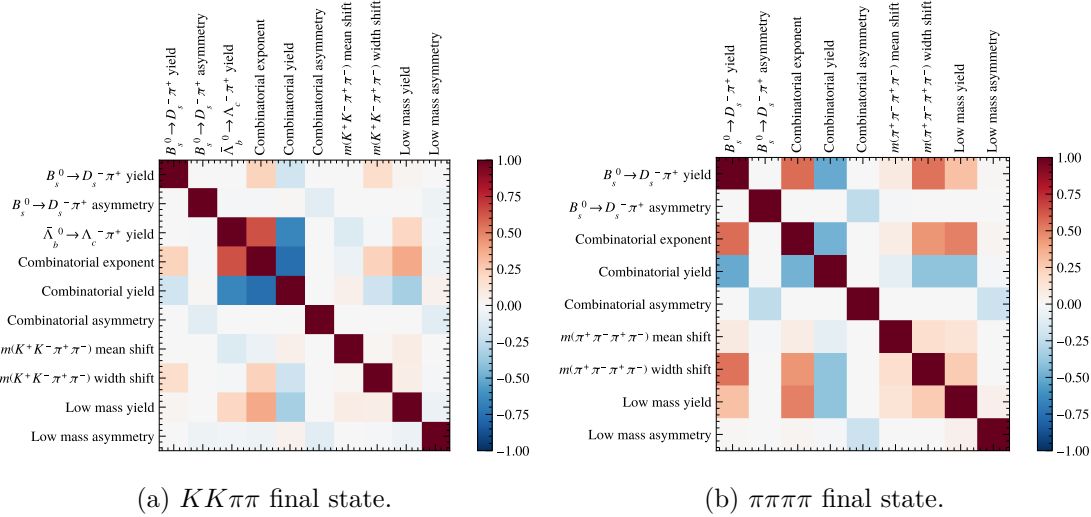
(a) $KK\pi\pi$ final state.(b) $\pi\pi\pi\pi$ final state.

Figure 4.12: Correlations between parameters of the nominal fits, as listed in Table 4.7 and shown in Fig. 4.11.

prevent any such conclusions being drawn, with all asymmetry parameters consistent with 0.

The (anti)correlations between parameters are as expected. A_{raw} exhibits only a weak anticorrelation with the combinatorial background asymmetry, and a yet weaker anticorrelation with the low mass background asymmetry, both of which are anticipated. The lack of correlations between the asymmetries and yield parameters also indicates that the parameterisation of the N^\pm per Eq. 4.12 is effective.

In the $KK\pi\pi$ final state, the strongest (anti)correlations are those between the combinatorial exponent and the yields of combinatorial and $A_b^0 \rightarrow \bar{A}_c^- \pi^+$ backgrounds. These are also anticipated, as changes in the shape of the combinatorial background via the combinatorial exponent induce changes in the combinatorial yield, which are then compensated for by the $A_b^0 \rightarrow \bar{A}_c^- \pi^+$ yield as the $A_b^0 \rightarrow \bar{A}_c^- \pi^+$ background is broad and lies between the $B_s^0 \rightarrow D_s^- \pi^+$ peak and the upper $m(KK\pi\pi)$ sideband.

In the $\pi\pi\pi\pi$ final state, the strongest (anti)correlations are largely induced by the combinatorial exponent: as the low mass background is effectively the only background lying to the left of the $B_s^0 \rightarrow D_s^- \pi^+$ peak, the events at lower $m(\pi\pi\pi\pi)$ can be comfortably modelled by either the low mass or combinatorial background components, hence an increase in the combinatorial exponent. This effect also introduces an instability in the width of the low mass background, reflected in the strong correlation between the combinatorial exponent and δ_σ . As δ_σ varies, the area beneath the $B_s^0 \rightarrow D_s^- \pi^+$, and consequently the $B_s^0 \rightarrow D_s^- \pi^+$ yield, varies, resulting in a similarly strong correlation between $N_{B_s^0 \rightarrow D_s^- \pi^+}$ and δ_σ , and the combinatorial exponent/yield.

4.3.3 Fit validation and pseudoexperiments

The nominal invariant mass fits must be validated thoroughly to ensure that the values of A_{raw} returned are reliable and robust. This analysis makes extensive use of pseudoexperiments to validate the nominal fit, in which many pseudo-datasets are generated according to the nominal fit model, on which the nominal fit is then performed, comparing the resulting parameters with those of the nominal fits [113]. Of particular interest are the errors on the parameters in the pseudoexperiment fits and the pulls between the values of each parameter in the nominal and pseudoexperiment fits:

$$\mathcal{P}_x = \frac{x_{\text{pseu.}} - x_{\text{nom.}}}{\sigma_{x,\text{pseu.}}}, \quad (4.16)$$

for a parameter x measured to be $x_{\text{pseu.}}$ and $x_{\text{nom.}}$ in a given pseudoexperiment fit and the nominal fit, respectively, and $\sigma_{x,\text{pseu.}}$ is the uncertainty on $x_{\text{pseu.}}$. If the fit is unbiased, then each value of $x_{\text{pseu.}}$ can be considered to be drawn from some probability distribution centred on $x_{\text{nom.}}$ with a variance estimated by $\sigma_{x,\text{pseu.}}^2$. According to the central limit theorem, if $x_{\text{pseu.}}$ is sampled sufficiently many times, then the distribution of sampled values will tend towards a Gaussian of width $\sigma_{x,\text{pseu.}}$ and mean $x_{\text{nom.}}$. In other words, an unbiased fit should return pulls which are normally distributed. By fitting a Gaussian distribution to the distribution of pulls, the mean and width of said distribution can be compared to their anticipated values of 0 and 1, respectively. Similarly, the distribution of $\sigma_{x,\text{pseu.}}$ itself should follow a Gaussian distribution centred on the statistical uncertainty on x .

To study the sensitivity and stability of the nominal mass fit, 1000 pseudoexperiments were performed in each final state, with 852 of the fits converging in $KK\pi\pi$ and 932 converging in $\pi\pi\pi\pi$. The resulting pulls and pseudoexperiment fit errors for A_{raw} , to which Gaussian distributions were fit, are shown in Fig. 4.13.

The Gaussian distributions describe the pulls and errors well in both final states, returning $\mu = 0.001 \pm 0.034$ and $\sigma = 1.013 \pm 0.025$ in $KK\pi\pi$ and $\mu = 0.013 \pm 0.034$ and $\sigma = 1.034 \pm 0.024$ in $\pi\pi\pi\pi$, consistent with those of a normal distribution. The fits to the errors on A_{raw} in the pseudoexperiments return statistical sensitivities of

$$\sigma_{\text{stat.}}(KK\pi\pi) = (6.20557 \pm 0.00077) \times 10^{-3},$$

$$\sigma_{\text{stat.}}(\pi\pi\pi\pi) = (1.60970 \pm 0.00080) \times 10^{-2}.$$

Gaussian distributions were also fit to the distributions of pulls of the other floating parameters in the nominal mass fit. The Gaussian means and widths (subtracted by 1 to be centred on 0) resulting from these fits are shown in Fig. 4.14.

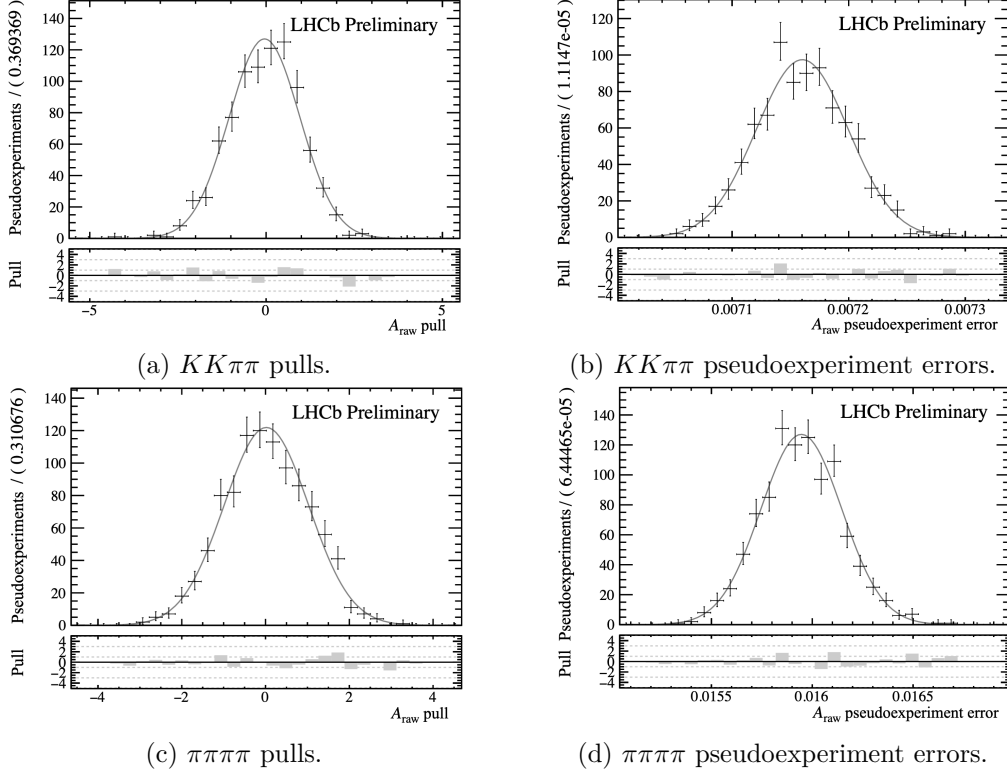


Figure 4.13: Pulls and pseudoexperiment errors on A_{raw} from the nominal pseudoexperiment study. Gaussian distributions fit to each of these are overlaid.

Generally, the pulls on these parameters are consistent with being normally distributed. For those parameters for which this is not the case, the scales of biases are very small, with all biases (displacements of the pulls mean from 0) below 0.15.

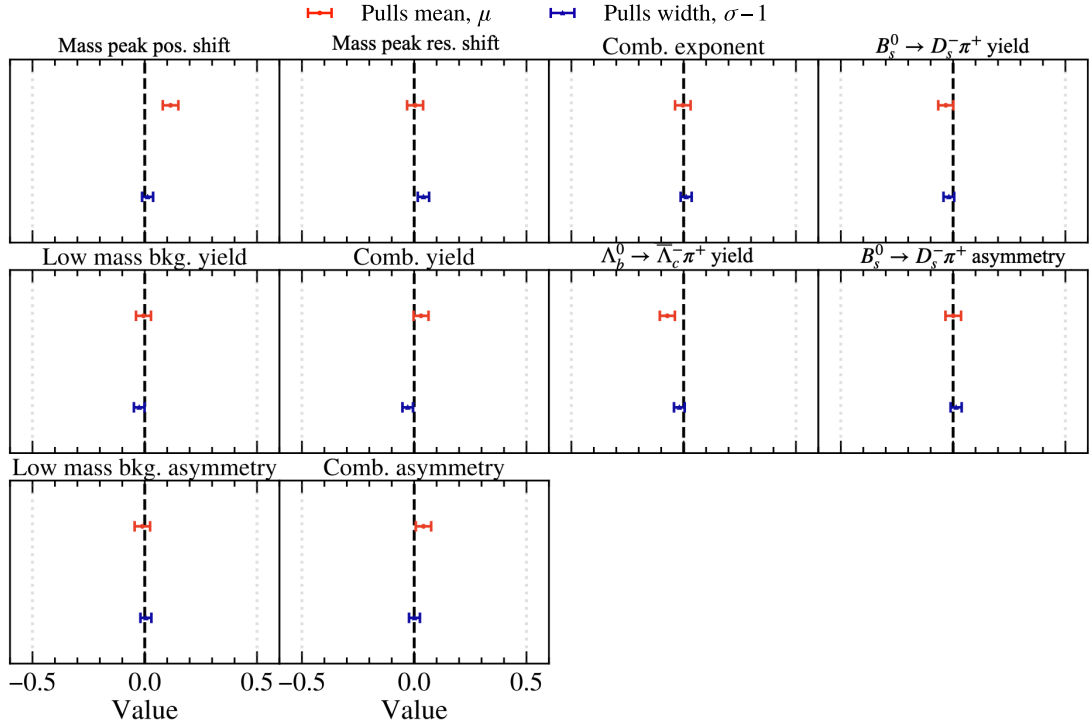
Further to the pseudoexperiment studies described above, the fit was validated by comparison to a charge-agnostic fit, *i.e.*, in which no distinction is made on the final-state charge and hence no asymmetry parameters are incorporated. This was performed on the same dataset as the nominal fit, with the resulting distribution shown in Fig. 4.15, with corresponding parameters listed in Table 4.8.

Table 4.8: Parameters resulting from the charge-agnostic fit, as shown in Fig. 4.15


Parameter	Fit value	
	$KK\pi\pi$ final state	$\pi\pi\pi\pi$ final state
$B_s^0 \rightarrow D_s^- \pi^+$ yield	28940 ± 180	7270 ± 150
Combinatorial yield	8950 ± 180	37620 ± 250
Low mass yield	1512 ± 57	720 ± 68
$A_b^0 \rightarrow \bar{A}_c^- \pi^+$ yield	610 ± 110	—
Comb. exponent / $(\text{MeV}/c^2)^{-1}$	$(-8.9 \pm 3.0) \times 10^{-4}$	$(-1.186 \pm 0.057) \times 10^{-3}$
Mass value shift / MeV/c^2	-0.63 ± 0.12	-0.93 ± 0.36
Mass resolution shift / MeV/c^2	8.31 ± 0.39	3.41 ± 0.82

As expected, all parameters common between the nominal and charge-agnostic fits are in

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays



(a) $KK\pi\pi$ final state.



(b) $\pi\pi\pi\pi$ final state.

Figure 4.14: Widths and means of Gaussian distributions as fit to the pulls of each floating parameter in the nominal fit.

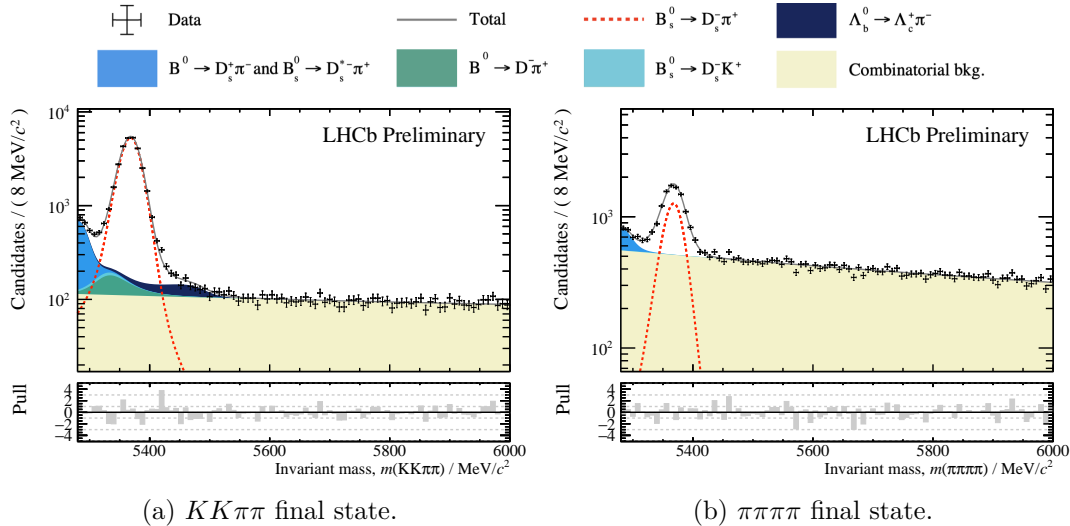


Figure 4.15: Charge-agnostic equivalent to the nominal invariant mass fits shown in Fig. 4.11. Corresponding parameters are listed in Table 4.8.

good agreement, with near-identical central values and errors. This, as for the nominal fit parameter correlations, stands to indicate that the parameterisation of N^\pm is effective, as the asymmetry parameters should, and indeed do, describe any differences between the two fit configurations.

Two other effects can also be probed through the use of charge-agnostic fits: the fits described thus far employ a bespoke Python- and RooFit-based fitting framework, and is therefore not validated by use in prior analyses; the modelling of the combinatorial and low mass backgrounds in the fit do not use the information of the $hh\pi\pi$ invariant mass distribution below $5100 \text{ MeV}/c^2$. To address these, the charge-agnostic invariant mass fit described above is first performed in the C++- and RooFit-based `B2DXFitters` framework, using the same shape parameters as obtained in Section 4.3.1. The $hh\pi\pi$ invariant mass window is then extended at the lower end to $5100 \text{ MeV}/c^2$ and descriptions of the $B_s^0 \rightarrow D_s^- \rho^+$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$ included, parameterised by the `RooHORNSdini` [114] and sum of a `RooHORNSdini` and `RooHILLdini` [114], respectively, with the shapes obtained by fits to MC. As the latter fit has the additional information from the full structures of the partially reconstructed backgrounds, their correlation with the combinatorial background is significantly reduced, highlighting any bias induced by the combinatorial and low mass modelling considerations in Section 4.3.1.

Both of these fits return yields consistent with those of the nominal invariant mass fit for both final states. In the fit over the extended $m(hh\pi\pi)$ range, shown in Fig. 4.16, signal yields of 28930 ± 190 and 7380 ± 150 are obtained for the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively. The agreement of these yields with those of the nominal invariant mass fit of 0.10σ in $KK\pi\pi$ and 0.06σ in $\pi\pi\pi\pi$.

4.3.4 Asymmetry contributions from background components

Provided that $B_s^0 \rightarrow D_s^- \pi^+$ candidates are adequately separated from the background components in $m(hh\pi\pi)$, asymmetries in background should not contribute directly to $\langle A_{\text{untagged}}^s \rangle$. This is in contrast to similar analyses in semi-leptonic b -hadron decays, *e.g.*, Ref. [115], where an additional asymmetry term is introduced to account for asymmetries in background components. However, the presence of any asymmetry in the background

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

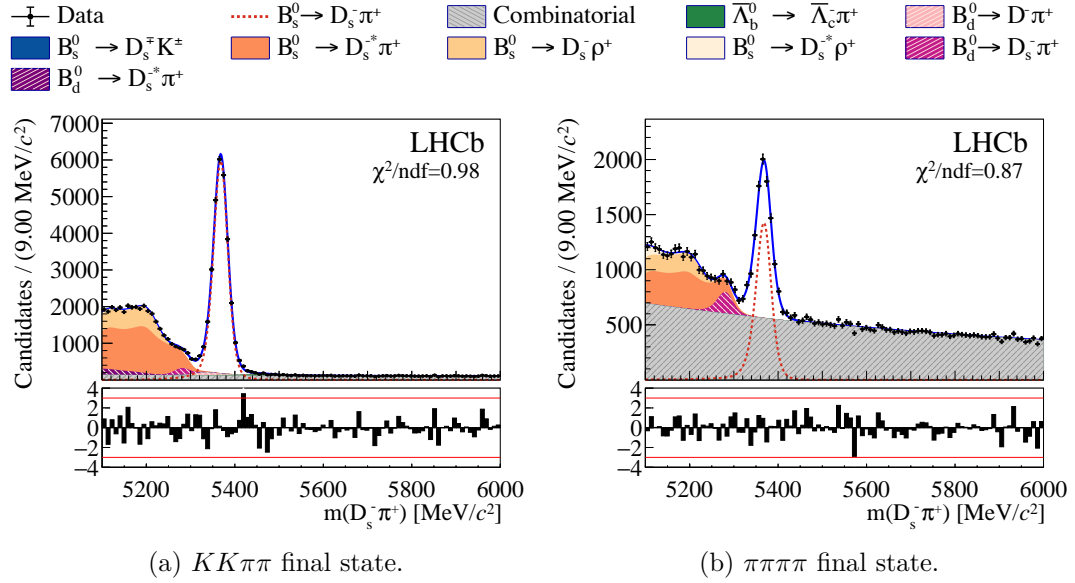
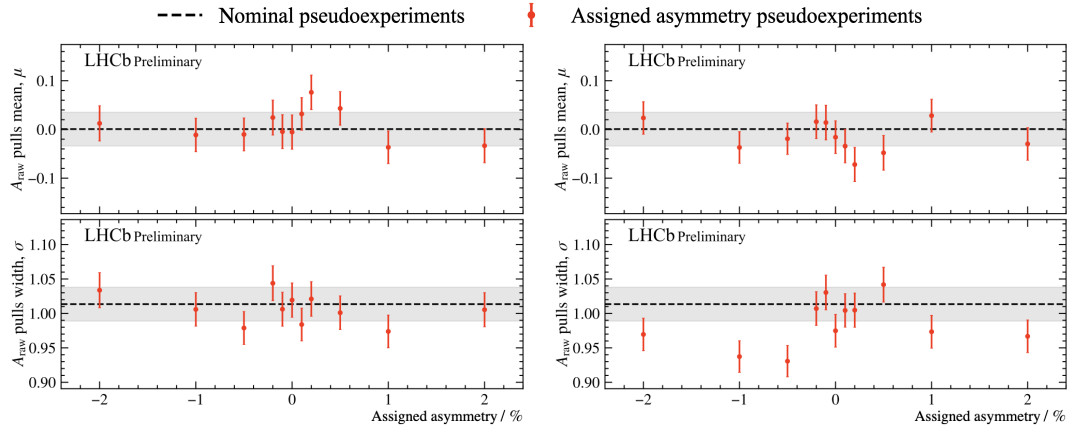


Figure 4.16: Charge-agnostic fit to $m(hh\pi\pi)$ using `B2DXFitters` with the lower limit of the mass window extended to $5100 \text{ MeV}/c^2$.



contributions, particularly those for which the asymmetry is fixed to 0, could induce a bias in the determination of A_{raw} .

To evaluate this bias, the nominal pseudoexperiment studies of Section 4.3.3 is repeated; however, the MC samples are generated with an asymmetry assigned to a given background component. In the both final states, asymmetries are assigned in the combinatorial, low mass and $B_s^0 \rightarrow D_s^- K^+$ components; $B^0 \rightarrow D^- \pi^+$ and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ are also included in the $KK\pi\pi$ final state. For each component, asymmetries of 0, $\pm 0.1\%$, $\pm 0.2\%$, $\pm 0.5\%$, $\pm 1\%$, $\pm 2\%$ were assigned. As in Section 4.3.3, Gaussian distributions are fit to the pulls on A_{raw} in each component-asymmetry scenario, with the resulting means/widths compared to their equivalents in the nominal studies. If a background asymmetry induces a bias in A_{raw} , then the mean of the A_{raw} pulls should be clearly dependent on the assigned asymmetry. The width of the A_{raw} pulls distribution is only expected to be modified by a background asymmetry through differences in the fit stability. These studies were performed across 1000 pseudoexperiments in each component-asymmetry scenario, with the resulting A_{raw} pulls means and widths shown for the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states in Figs. 4.17 & 4.18, respectively.

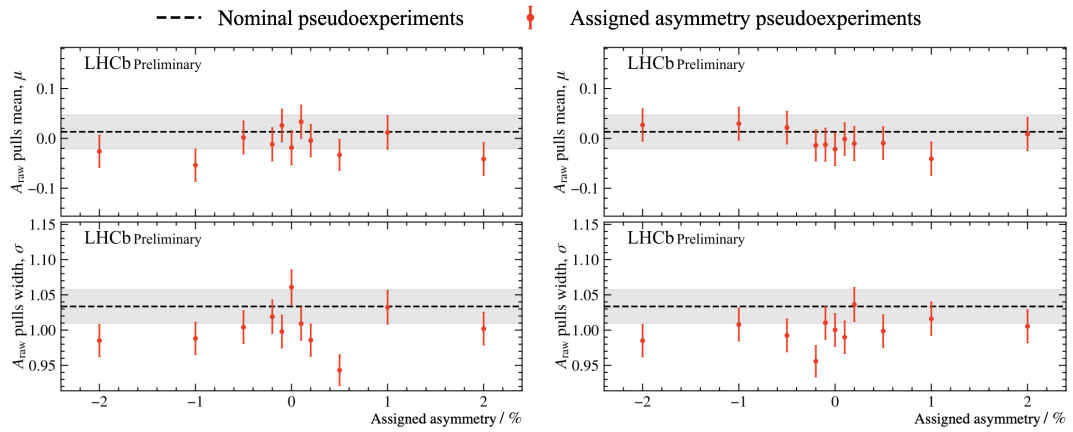
In the $KK\pi\pi$ final state, the means and widths of most component-asymmetry scenarios were consistent with those of the nominal fit, *i.e.*, that the nominal fits are not biased by a background asymmetry in any of the components up to the scale of 5%. Additionally, the small variations observed in these parameters appear generally random, rather than exhibiting a significant trend with the assigned asymmetry. In the low mass background asymmetry scenarios, for larger asymmetries ($|A| > 0.5\%$), the distributions of pulls are slightly narrower. This is understood to arise from the changes that such a background asymmetry introduces to the correlations between the $B_s^0 \rightarrow D_s^- \pi^+$, combinatorial and low mass background. In the $B^0 \rightarrow D^- \pi^+$ scenarios, the slight negative bias for negative asymmetries and slight positive bias for positive asymmetries could be a sign as a dependence of the bias on the $B^0 \rightarrow D^- \pi^+$ background asymmetry, arising from the fact that the $B^0 \rightarrow D^- \pi^+$ asymmetry is fixed to 0 in the nominal invariant mass fits. However, the scale of this bias is small, and the systematic uncertainty assigned to describe the




(a) Combinatorial background asymmetries assigned. (b) Low mass background asymmetries assigned.

(c) $B_s^0 \rightarrow D_s^- K^+$ asymmetries assigned.(d) $B^0 \rightarrow D^- \pi^+$ asymmetries assigned.(e) $A_b^0 \rightarrow \bar{A}_c^- \pi^+$ asymmetries assigned.Figure 4.17: Gaussian means and widths, as fit to $A_{\text{raw}}(KK\pi\pi)$ for scenarios of assigned asymmetries in each background component, with those of the nominal fit, per Fig. 4.13, overlaid (grey/black).

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays



(a) Combinatorial background asymmetries assigned. (b) Low mass background asymmetries assigned.



(c) $B_s^0 \rightarrow D_s^- K^+$ asymmetries assigned.

Figure 4.18: Gaussian means and widths, as fit to $A_{\text{raw}}(\pi\pi\pi\pi)$ for scenarios of assigned asymmetries in each background component, with those of the nominal fit, per Fig. 4.13, overlaid (grey/black).

fixing of the asymmetry of the $B^0 \rightarrow D^- \pi^+$ component to 0 (discussed in Section 4.6.1) should describe precisely this behaviour.

The same is true in the $\pi\pi\pi\pi$ final state, where effectively all means and widths are consistent with their nominal equivalents. All variations appear to be random, exhibiting no dependence on the assigned background asymmetries.

As the A_{raw} pull means and widths of the component-asymmetry scenarios are all consistent or close to consistent with those of the nominal pseudoexperiment studies in Section 4.3.3, and since no significant dependence on the assigned background asymmetries are observed, it is concluded that the effect of any background asymmetries on the nominal fit is significantly below the statistical sensitivity to A_{raw} and hence can be safely neglected. Additionally, any biases which could arise in the background components with asymmetries fixed to 0 in the nominal fit are accounted for in the corresponding systematics discussed in Section 4.6.1.

4.4 Detection asymmetry, $A_{\text{det.}}$

Assuming that contributions to the detection asymmetry can be factorised, $A_{\text{det.}}$ can be decomposed as

$$A_{\text{det.}} = A_{\text{det.}}^{\text{sel.}} + A_{\text{det.}}^{\text{ProbNN}} + A_{\text{det.}}^{\text{L0}} + A_{\text{det.}}^{\text{track}}, \quad (4.18)$$

where each term describes the contribution from asymmetries in the selection⁵, particle identification, L0 trigger and track reconstruction, respectively. This assumption only holds if there are no significant local asymmetries in the distributions in which these components are determined, namely the component momenta of the B_s^0 , D_s^- , and final-state particles, and quantities derived directly from these, *e.g.*, p_{T} and η . To test this, charge asymmetries of the MC/data samples were investigated; no significant local asymmetries were identified, the assumption of factorisation was deemed valid.

Each of the components in Eq. 4.18 can be evaluated by considering the corresponding efficiency in each charge-conjugate final state. The methods by which the efficiencies of these contributions are estimated, and their resulting values, are discussed in Section 4.2. As in Ref. [116], the detection asymmetry can be obtained from the charge-specific efficiencies by incorporating the efficiencies into their respective yields, *i.e.*,

$$N^\pm \rightarrow N'^\pm = N^\pm \varepsilon^\pm. \quad (4.19)$$

Under this transformation, the asymmetry (as in Eq. 4.6) is transformed to

$$A \rightarrow A' = \frac{N^- \varepsilon^- - N^+ \varepsilon^+}{N^- \varepsilon^- + N^+ \varepsilon^+}. \quad (4.20)$$

This can alternatively be expressed in terms of the sum, $\varepsilon_\Sigma = \varepsilon^- + \varepsilon^+$, and difference, $\varepsilon_\Delta = \varepsilon^- - \varepsilon^+$, of the charge-specific efficiencies:

$$A' = \frac{N^- (\varepsilon_\Sigma + \varepsilon_\Delta) - N^+ (\varepsilon_\Sigma - \varepsilon_\Delta)}{N^- (\varepsilon_\Sigma + \varepsilon_\Delta) + N^+ (\varepsilon_\Sigma - \varepsilon_\Delta)}, \quad (4.21)$$

which can be rearranged for a numerator and denominator expressed in factors of $\varepsilon_{\Sigma,\Delta}$,

$$A' = \frac{\varepsilon_\Sigma (N^- - N^+) + \varepsilon_\Delta (N^- + N^+)}{\varepsilon_\Sigma (N^- + N^+) + \varepsilon_\Delta (N^- - N^+)}. \quad (4.22)$$

⁵In this section, the selection efficiency/asymmetry refers solely to those of the kinematic selection requirements, *e.g.*, the invariant mass requirement.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

Eliminating a factor of $\varepsilon_\Sigma (N^- + N^+)$ from both numerator and denominator, and substituting in A , this becomes

$$A' = \frac{A + \varepsilon_\Delta / \varepsilon_\Sigma}{1 + A \varepsilon_\Delta / \varepsilon_\Sigma} \quad (4.23)$$

For small asymmetries in the yields and efficiencies, this approximates to

$$A' \simeq A + \frac{\varepsilon^- - \varepsilon^+}{\varepsilon^- + \varepsilon^+}, \quad (4.24)$$

where the second term is the contribution to the asymmetry of the efficiencies, from which each $A_{\text{det.}}^i$ in Eq. 4.18 is computed.

4.4.1 Contributions from kinematic selection efficiencies

The kinematic selection asymmetry, $A_{\text{det.}}^{\text{sel.}}$, describes the asymmetry between the efficiencies of three offline selections: the invariant mass windows, $B^0 \rightarrow D^{*-} \pi^+$ veto and HCAL separation requirement, in each charge-conjugate final state. These efficiencies are computed for each year and magnet polarity in $B_s^0 \rightarrow D_s^- \pi^+$ samples and are listed in Table 4.3 in Section 4.2.1, and combined as a luminosity-weighted mean.

The corresponding asymmetries, $A_{\text{det.}}^{\text{sel.}}$ in these efficiencies are listed for the efficiencies in the magnet up and down configurations averaged across data-taking years, and averaged across both polarities and all data-taking years in Table 4.9. As anticipated, the asym-

Table 4.9: Charge-specific selection efficiencies and corresponding asymmetries from $B_s^0 \rightarrow D_s^- \pi^+$ MC samples.

Final state	Magnet	Efficiencies / %		Asymmetry / %
		$D_s^+ \pi^-$	$D_s^- \pi^+$	
$KK\pi\pi$	Up	24.017 ± 0.030	24.035 ± 0.030	-0.037 ± 0.089
	Down	24.117 ± 0.030	24.069 ± 0.030	0.099 ± 0.088
	Both	24.067 ± 0.021	24.052 ± 0.021	0.031 ± 0.063
$\pi\pi\pi\pi$	Up	23.94 ± 0.12	23.97 ± 0.12	-0.05 ± 0.34
	Down	23.96 ± 0.11	23.91 ± 0.11	0.10 ± 0.33
	Both	23.949 ± 0.081	23.937 ± 0.081	0.03 ± 0.24

metry arising from the kinematic selection is negligibly small, with the largest tension of 1.1σ in the 2016-18 combinations arising in the magnet down configuration and $KK\pi\pi$ final state. A systematic uncertainty associated with the factorisation of PID and L0 efficiencies from $A_{\text{det.}}^{\text{sel.}}$ is described in Section 4.6.2.

4.4.2 Contribution from PID efficiencies

The term $A_{\text{det.}}^{\text{ProbNN}}$ arises from the charge-specific efficiencies of particle identification requirements: the cuts on ProbNN variables, `hasRich` and `isMuon`. These are, per Section 4.2.2, in the resonant and non-resonant regions of the $D_s^+ \rightarrow h^+ h^- \pi^+$ phase space, using calibration samples of $D^{*+} \rightarrow D^0 \pi^+$ with $D^0 \rightarrow K^- \pi^+$ via `PIDCalib2`. The efficiencies are combined per-year and magnet polarity as weighted mean where the relative signal yields in the (non-)resonant regions are used as weights, which are further combined by year/magnet polarity as luminosity-weighted means. $A_{\text{det.}}^{\text{ProbNN}}$ is then determined from these efficiencies according to Eq. 4.24.

Systematic uncertainties accounting for the granularity of the binning scheme applied to the calibration samples and the choice of `ProbNNk` and `ProbNNpi` selections are discussed in Section 4.6.2.

$A_{\text{det.}}^{\text{ProbNN}}$ is determined for the 2016-18 combination in magnet up/down and combined cases, listed in Table 4.10. As for $A_{\text{det.}}^{\text{sel.}}$, the values of $A_{\text{det.}}^{\text{ProbNN}}$ are negligibly small, with all values averaging between resonant and non-resonant regions consistent with 0.

Table 4.10: Charge-specific PID asymmetries in $B_s^0 \rightarrow D_s^- \pi^+$, in the (non-)resonant $D_s^+ \rightarrow h^+ h^- \pi^+$ phase space regions, and their luminosity-weighted means.

Final state	Magnet	Asymmetry / %		
		Non-resonant	Resonant	Average
$KK\pi\pi$	Up	0.07 ± 0.14	0.09 ± 0.30	0.07 ± 0.13
	Down	-0.01 ± 0.13	-0.16 ± 0.30	-0.03 ± 0.12
	Both	0.028 ± 0.096	-0.036 ± 0.213	0.021 ± 0.088
$\pi\pi\pi\pi$	Up	0.29 ± 0.76	0.35 ± 0.60	0.31 ± 0.49
	Down	-0.05 ± 0.74	-0.23 ± 0.59	-0.14 ± 0.48
	Both	0.12 ± 0.53	0.06 ± 0.42	0.08 ± 0.34

4.4.3 Contribution from L0 trigger efficiencies

Similarly to the PID detection asymmetry, the asymmetry in charge-specific L0 trigger efficiencies can be computed from those efficiencies discussed and estimated in Section 4.2.3 using calibration samples of $D^{*+} \rightarrow D^0 \pi^+$ with $D^0 \rightarrow K^- \pi^+$ via `L0HadronTables`.

The asymmetries $A_{\text{det.}}^{\text{L0}}$ are computed per Eq. 4.24, with the statistical uncertainty from the $B_s^0 \rightarrow D_s^- \pi^+$ MC samples taken as the statistical uncertainty. A systematic uncertainty for the calibration sample statistics is discussed in Section 4.6.2. The charge-specific efficiencies and corresponding asymmetries are listed in Table 4.11.

Table 4.11: Charge-specific L0 selection efficiencies and corresponding asymmetries from $B_s^0 \rightarrow D_s^- \pi^+$ MC samples.

Final state	Magnet	Efficiencies / %		Asymmetry / %
		$D_s^+ \pi^-$	$D_s^- \pi^+$	
$KK\pi\pi$	Up	39.024 ± 0.044	38.961 ± 0.044	-0.081 ± 0.079
	Down	39.361 ± 0.044	39.161 ± 0.043	-0.255 ± 0.078
	Both	39.193 ± 0.031	39.061 ± 0.031	-0.169 ± 0.056
$\pi\pi\pi\pi$	Up	41.32 ± 0.18	41.48 ± 0.18	0.19 ± 0.30
	Down	41.76 ± 0.18	41.60 ± 0.17	-0.19 ± 0.29
	Both	41.54 ± 0.12	41.54 ± 0.12	-0.00 ± 0.21

In the $KK\pi\pi$ final state, a significant non-zero contribution to the detection asymmetry is measured, both for the individual magnet polarities and in their combination. No such significant contribution is measured in the $\pi\pi\pi\pi$ final state.

4.4.4 Tracking asymmetry

Finally, the asymmetry induced by charge-dependent differences in track reconstruction efficiencies is evaluated. This involves determining the efficiencies to reconstruct tracks or pairs of tracks, which depends both on the efficiency to detect particles in the VELO, TT and T1-T3, and the efficiencies to reconstruct long tracks from these. Pions are not considered to exhibit a significant charge-dependence in their interactions with the detector, *i.e.*, their nuclear cross-sections are not significantly charge-dependent. On the other hand, the nuclear cross-sections of kaons exhibit a strong charge-dependence: the \bar{u} in a K^- is available to annihilate with the u -quarks in the protons and neutrons of nuclei, whereas neither of the quarks in a K^+ is present in either nucleon, except via their sea quarks. This effect can induce CP asymmetries up to the scale of $\mathcal{O}(10^{-2})$, as determined from the cross-sections of kaons with deuterium [5]. Earlier LHCb measurements in Run 1 established that the asymmetry in the reconstruction of a $K^\pm \pi^\mp$ pair could be determined more accurately and reliably than through the separate determination for a K^\pm and π^\mp individually [117]. Whilst the single-pion asymmetry can be determined reliably by estimating the charge-specific tracking efficiencies with $J/\psi \rightarrow \mu^+ \mu^-$ decays per Section 4.2 and applying Eq. 4.24, the charge-dependent nuclear interactions of kaons prevent the same method from being applied for an equivalent single-kaon tracking asymmetry. Instead, this analysis evaluates the single-kaon tracking asymmetry by measuring the $K^\pm \pi^\mp$ and single-pion tracking asymmetries, and taking their difference.

Considering the final states of the analysis, looking first to the $KK\pi\pi$ final state, the total tracking asymmetry can be constructed from the tracking asymmetries of three components: a $K^- \pi^+$ pair and single K^+ (both originating from $D_s^+ \rightarrow K^+ K^- \pi^+$), and a single π^- (the associated pion). The $\pi\pi\pi\pi$ final state is much simpler to compute, taking the sum of single-pion asymmetries, whilst minding their relative charges, *i.e.*, that the π^+ and π^- tracking symmetries have a relative sign difference. These are written out as

$$A_{\text{det.}}^{\text{track}}(KK\pi\pi) = A_{\text{det.}}^{\text{track}}(K^+) + A_{\text{det.}}^{\text{track}}(\pi^-) + A_{\text{det.}}^{\text{track}}(K^- \pi^+), \quad (4.25a)$$

$$A_{\text{det.}}^{\text{track}}(\pi\pi\pi\pi) = A_{\text{det.}}^{\text{track}}(\pi_{B_s^0}^-) + A_{\text{det.}}^{\text{track}}(\pi_{D_s^+,1}^+) + A_{\text{det.}}^{\text{track}}(\pi_{D_s^+}^-) + A_{\text{det.}}^{\text{track}}(\pi_{D_s^+,2}^+). \quad (4.25b)$$

As both final states are charge-symmetric, the asymmetries between final-state particles should cancel significantly within kinematic differences (especially if the K^+ and K^- originate from a ϕ), *i.e.*, a smaller cancellation is anticipated between the π^+ of the D_s^+ and the associated pion.

The two-track reconstruction asymmetry of the $K^- \pi^+$ pair, $A_{\text{det.}}^{\text{track}}(K^- \pi^+)$, from a D_s^+ decay has been studied in detail in previous analyses, with a standard method established for determining these, implemented in a Python-based common software tool, `KPiAsymmetryTool` [117]. Assuming CP conservation in Cabbibo-favoured charm decays, $A_{\text{det.}}^{\text{track}}(K^- \pi^+)$ can be determined from the raw asymmetries of $D^+ \rightarrow K^- \pi^+ \pi^+$ and $D^+ \rightarrow K_S^0 \pi^+$, corrected by the asymmetry of neutral kaons, $A(K^0)$, such that

$$A_{\text{det.}}^{\text{track}}(K^- \pi^+) = A_{\text{raw}}(D^+ \rightarrow K^- \pi^+ \pi^+) - A_{\text{raw}}(D^+ \rightarrow K_S^0 \pi^+) - A(K^0). \quad (4.26)$$

$A_{\text{raw}}(D^+ \rightarrow K^- \pi^+ \pi^+)$ and $A_{\text{raw}}(D^+ \rightarrow K_S^0 \pi^+)$ can be estimated from dedicated calibration samples, whilst $A(K^0)$ is taken from a bespoke simulation in Ref. [118] to be $A(K^0) = (0.054 \pm 0.014)\%$. The calibrations samples are sourced from HLT2 trigger lines saving to the Turbo steam: `Hlt2CharmHadDpToKmPipPip_ForKPiAsymTurbo` and

Hlt2CharmHadDpKSOPip_KSOLLAsymTurbo for $D^+ \rightarrow K^- \pi^+ \pi^+$ and $D^+ \rightarrow K_S^0 \pi^+$, respectively.

To mitigate the impact of differing production and trigger asymmetries between the calibration channels, per-event weights are applied to the sample of $D^+ \rightarrow K_S^0 \pi^+$ decays such that the D^+ are kinematically similar between the two samples. The samples must also be weighted such that they are kinematically similar to the $D_s^+ \rightarrow K^+ K^- \pi^+$ from $B_s^0 \rightarrow D_s^- \pi^+$; per-event *sWeights* are extracted from the nominal invariant mass fit, and used to subtract background contributions from the sample of data, such that the *sWeighted* data samples represent only $B_s^0 \rightarrow D_s^- \pi^+$ contributions. The weighting procedure thus involves first weighting the $D^+ \rightarrow K^- \pi^+ \pi^+$ calibration sample to match the kinematics of the $D_s^+ \rightarrow K^+ K^- \pi^+$ in the *sWeighted* data samples of this analysis, weighting in (p, η) of the common K^- and (p_T, η) of the common π^+ . The $D^+ \rightarrow K_S^0 \pi^+$ calibration sample is then weighted to match the kinematics of the weighted $D^+ \rightarrow K^- \pi^+ \pi^+$ calibration sample, weighting in p_T of the common π^+ and (p_T, η) of the common D^+ . Finally, the weighted $D^+ \rightarrow K_S^0 \pi^+$ calibration sample to the $D^+ \rightarrow K^- \pi^+ \pi^+$ calibration sample in ϕ , the angle between the particle trajectory in the x - y plane and the x -axis, of the common π^+ .

Two methods were considered for performing the weighting described above. A machine learning (ML) classifier, namely the boosted decision tree-based **GBReweigher** from the **hep_ml** Python package [119], is trained to predict the per-event weights necessary to perform a given kinematic reweighting. Alternatively, a binned weighting, namely **BinsReweigher** also implemented in **hep_ml**, in which the weights are taken simply as the ratio between histograms of the target and original distribution, with an option of applying a Gaussian filter to include information from neighbouring bins. The former method is applied here, with the latter used in the determination of a systematic uncertainty due to the reweighting (see Section 4.6.2). An example of the first stage of the weighting (weighting the $D^+ \rightarrow K^- \pi^+ \pi^+$ calibration sample to the *sWeighted* data sample from this analysis), for data taken in 2016 with the magnet up configuration, is shown in Fig. 4.19.

With the calibration samples weighted, the charge-specific yields, and subsequent values of A_{raw} , in each D^+ decay mode are obtained by fits to the $K^- \pi^+ \pi^+ / K_S^0 \pi^+$ invariant masses. The signal contributions are modelled by a DSCB, a PDF formed of the sum of three distinct Gaussian distributions, convolved with a power-law function to model the radiative tails of the D^+ mass peaks. The small combinatorial background contributions, the only backgrounds considered present in these samples, are modelled by single exponential functions. Fits were performed to the samples in each data-taking year and magnet polarity, in bins of p and η of the K^- in $D^+ \rightarrow K^- \pi^+ \pi^+$, with the procedure carried out for both the $K^- \pi^+$ and $K^+ \pi^-$ pairs of the $K^+ K^- \pi^+ \pi^-$ final state.

The single-pion tracking asymmetry is much simpler to determine, applying Eq. 4.24 to the charge-specific single-pion efficiencies in bins of p_T and η , estimated from $J/\psi \rightarrow \mu^+ \mu^-$ calibration samples via **TrackCalib** in Section 4.2.4. Two systematic uncertainties are assigned to these asymmetries: one arising from the choice of signal model in the fits to $m(\mu^+ \mu^-)$, the other from the granularity of the kinematic binning scheme. Both of these are discussed in Section 4.6.2.

Finally, the single-kaon asymmetry, namely that of the K^+ in the $K^+ K^- \pi^+ \pi^-$ final state, is determined by evaluating the two-track reconstruction asymmetry of the $K^+ \pi^+$ pair and subtracting the single-pion asymmetry of the π^+ . The single-pion asymmetry is evaluated by dividing the $B_s^0 \rightarrow D_s^- \pi^+$ MC sample into bins of the K^+ p_T and η and

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

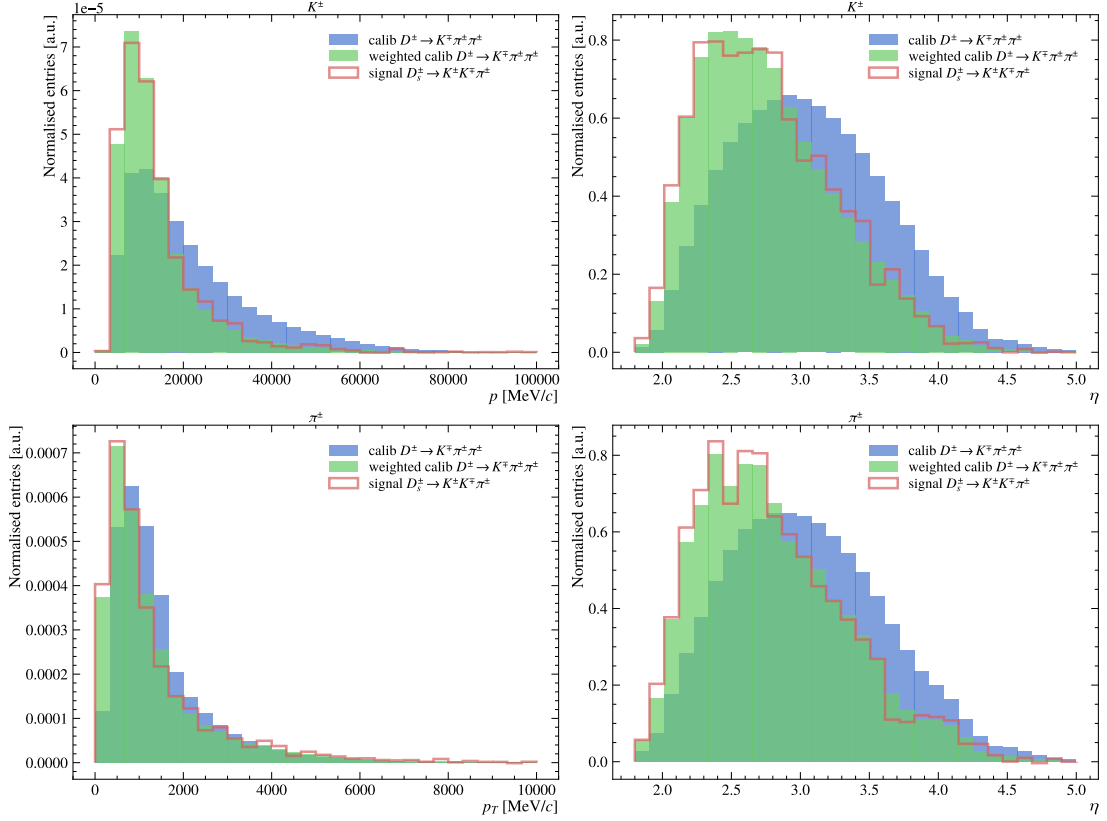


Figure 4.19: Weighting of the $D^+ \rightarrow K^- \pi^+ \pi^+$ calibration sample to the *sWeighted* data sample from this analysis, for data taken in 2016 with the magnet up configuration.

evaluating the average pion asymmetry in each bin from the single-pion efficiencies, such that a map of the asymmetry on π^+ in bins of the kaon kinematic bins is constructed.

The approach described above was applied for each year and magnet polarity to obtain values of $A_{\text{det.}}^{\text{track}}(hh\pi\pi)$. These values were then combined by year/magnet polarity according to their luminosity. The values of $A_{\text{det.}}^{\text{track}}(hh\pi\pi)$ and their constituent parts, combined per year and computed in the up, down and combination of magnet polarities, are listed in Tables 4.12 & 4.13 for the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively.

Table 4.12: Tracking detection asymmetry in $KK\pi\pi$ and its components.

Magnet polarity	Asymmetry / %			
	K^+	$(K^- \pi^+)_{D_s^+}$	$\pi_{B_s^0}^-$	$KK\pi\pi$
Up	1.10 ± 0.08	-1.08 ± 0.09	-0.08 ± 0.02	-0.06 ± 0.09
Down	1.04 ± 0.08	-1.01 ± 0.09	0.02 ± 0.02	0.05 ± 0.09
Both	1.07 ± 0.06	-1.04 ± 0.07	-0.03 ± 0.01	-0.00 ± 0.06

Two further effects must still be considered: the crossing angle of the proton beams and the hadronic interactions with which the muons used in determining the single-pion efficiencies/asymmetries differ from those of pions. Both of these effects are neglected, with corresponding systematics assigned in Section 4.6.2.

Table 4.13: Tracking detection asymmetry in $\pi\pi\pi\pi$ and its components.

Magnet polarity	Asymmetry / %			
	$\pi_{D_s^+}^+$	$(\pi^-\pi^+)_{D_s^+}$	$\pi_{B_s^0}^-$	$\pi\pi\pi\pi$
Up	-0.02 ± 0.04	0.00 ± 0.00	-0.06 ± 0.02	-0.08 ± 0.04
Down	0.04 ± 0.03	0.00 ± 0.00	0.01 ± 0.02	0.06 ± 0.03
Both	0.01 ± 0.02	0.00 ± 0.00	-0.02 ± 0.01	-0.01 ± 0.02

4.4.5 Summary of $A_{\text{det.}}$

The components of the detection asymmetry discussed above are combined according to Eq. 4.18. As the statistical uncertainties on $A_{\text{det.}}^{\text{sel.}}$, $A_{\text{det.}}^{\text{ProbNN}}$ and $A_{\text{det.}}^{L0}$ arise from exactly the same MC samples and are thus 100% correlated, the largest statistical uncertainty is taken of the three, that of $A_{\text{det.}}^{\text{ProbNN}}$ for both final states, and added in quadrature to the statistical uncertainty on $A_{\text{det.}}^{\text{track.}}$. The systematic uncertainties are computed as a quadrature sum over the systematic uncertainties on all four components.

The detection asymmetries in each final state are, per Eq. 4.18, determined to be

$$A_{\text{det.}}(KK\pi\pi) = (-0.12 \pm 0.11 \pm 0.09) \%,$$

$$A_{\text{det.}}(\pi\pi\pi\pi) = (0.00 \pm 0.21 \pm 0.02) \%.$$

These are both consistent with 0, indicating that the charge-symmetric final states provide a significant degree of cancellation in the contributions to the detection asymmetry.

4.5 Production asymmetry, $A'_{\text{prod.}}$

As discussed in Section 4.1, the production asymmetry, that is, the asymmetry inherent in the production of B_s^0 and \bar{B}_s^0 mesons in pp collisions, $A_{\text{prod.}}$, is washed out by the rapid B_s^0 oscillations. The parameter $A_{\text{prod.}}$ was measured in Ref. [29] to be -0.0031 ± 0.0032 . The dilution factor is a ratio of two integrals, both of which depend on the decay time acceptance, $\varepsilon(t)$. As the signal candidates selected in this measurement are sufficiently similar to those in Ref. [29], the decay time acceptance from Ref. [29] can be used. This acceptance is a cubic spline function, shown in Fig. 4.20.

Evaluating numerically the dilution factor given in Eq. 4.5, taking Γ_s and $\Delta\Gamma_s$ from Ref. [5] and Δm_s from Ref. [29], returns a ratio of $(7.320 \pm 0.017) \times 10^{-3}$. It is therefore safe to estimate that the effective production asymmetry, $A'_{\text{prod.}}$, the product of this factor and $A_{\text{prod.}}$, is, at its largest, $\mathcal{O}(10^{-5})$. As this is significantly below the sensitivity on A_{raw} achieved in this analysis, the production asymmetry term in Eq. 4.5 can be safely ignored.

4.6 Systematic uncertainties

Systematic uncertainties accounting for the effects discussed in the previous sections are assigned to A_{raw} and $A_{\text{det.}}$, summarised below in Tables 4.14 and 4.15, respectively.




Figure 4.20: Cubic spline function modelling the B_s^0 decay time acceptance, with knots and coefficients taken from Ref. [29].

4.6.1 Systematic effects in A_{raw}

The systematic effects on the determination of A_{raw} arise from how the invariant mass model is constructed and performed. Specifically, the systematic effects considered are the choice of signal model, the grouping $B^0 \rightarrow D_s^+ \pi^-$ and $B_s^0 \rightarrow D_s^{*-} \pi^+$ together using the $B^0 \rightarrow D_s^+ \pi^-$ model, and the fixing of parameters in the fit, namely the fraction between terms in the combinatorial model, the fractions of the signal yield from which the $B_s^0 \rightarrow D_s^- K^+$ and $B^0 \rightarrow D^- \pi^+$ yields are obtained, and the asymmetries of the $B_s^0 \rightarrow D_s^- K^+$, $B^0 \rightarrow D^- \pi^+$ and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ components. A standard approach is applied to quantify these effects, wherein the invariant mass model is modified to account for a given effect and fit to the data, with the difference between the resulting values of A_{raw} and those of the nominal fit taken as the systematic uncertainty. To quantify the choice of signal model, an invariant mass fit is performed taking the extended Cruijff+Gaussian PDFs shown in Figs. A.3c & A.3d as the $B_s^0 \rightarrow D_s^- \pi^+$ signal model (replacing also the $B^0 \rightarrow D_s^+ \pi^-$ model which derives from that of $B_s^0 \rightarrow D_s^- \pi^+$). For the modelling of low mass background, the $B^0 \rightarrow D_s^+ \pi^-$ shape is replaced by Johnson S_U functions fit to $B_s^0 \rightarrow D_s^{*-} \pi^+$ MC samples. For the parameters fixed in the nominal mass fit stated above, the alternative mass fits are performed with each parameter floating subject to a Gaussian constraint. The mean/width of the Gaussian constraint are taken as the central value/error of the parameter for the combinatorial and yield fractions, whilst a mean of 0 and width of 5% are used for the asymmetry parameters.

Performing the invariant mass fits described above, the systematic uncertainties associated with A_{raw} are listed in Table 4.14, along with the total systematic uncertainties (the quadrature sums over each contribution) for each final state. The parameterisation of the signal and low mass background components contributed significantly in both final states. In $KK\pi\pi$, the fraction between combinatorial background model terms and the fixing of the $B^0 \rightarrow D^- \pi^+$ and $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^- \pi^+$ produce yet greater contributions.

The systematic uncertainties on the values of A_{raw} in each final state measured in the data samples covering all data-taking years and magnet polarities are determined to be

$$\sigma_{\text{syst.}}(KK\pi\pi) = 4.9 \times 10^{-4},$$

$$\sigma_{\text{syst.}}(\pi\pi\pi\pi) = 1.3 \times 10^{-4}.$$

Table 4.14: Summary of systematic uncertainties on A_{raw} .

Systematic	Final state	
	$KK\pi\pi$	$\pi\pi\pi\pi$
Signal parameterisation	1.4×10^{-4}	1.0×10^{-4}
Low mass parameterisation	1.4×10^{-4}	9.0×10^{-5}
Combinatorial fraction	2.4×10^{-4}	2.9×10^{-6}
$B_s^0 \rightarrow D_s^- K^+$ yield fraction	4.6×10^{-8}	4.1×10^{-7}
$B^0 \rightarrow D^- \pi^+$ yield fraction	2.1×10^{-6}	—
$B_s^0 \rightarrow D_s^- K^+$ asymmetry	1.4×10^{-5}	4.7×10^{-7}
$B^0 \rightarrow D^- \pi^+$ asymmetry	3.6×10^{-4}	—
$A_b^0 \rightarrow \bar{A}_c^- \pi^+$ asymmetry	1.1×10^{-4}	—
Total	4.9×10^{-4}	1.3×10^{-4}

4.6.2 Systematic effects in $A_{\text{det.}}$

The systematic effects affecting $A_{\text{det.}}$, which are largely described in Section 4.4, are intertwined with the methods applied to determine each component of $A_{\text{det.}}$ (per Eq. 4.18).

For $A_{\text{det.}}^{\text{sel.}}$, systematic uncertainties describing the effect of neglecting/incorporating the PID and L0 requirements are assigned as half of the difference between the values of $A_{\text{det.}}^{\text{sel.}}$ without/with the corresponding efficiencies applied as per-event weights. This factor of a half, applied also for other systematics, is assigned to account for methods of estimating systematic uncertainties which consider more extreme variations than reasonably would be expected in an analysis.

Two effects are considered for $A_{\text{det.}}^{\text{ProbNN}}$: the first arising from the choice of binning scheme, the second from the choice of the cuts on ProbNN. The former is quantified as half of the difference between the value of $A_{\text{det.}}^{\text{ProbNN}}$ obtained with the standard binning scheme and the value obtained when the widths of the bins in both dimensions are halved. The latter is quantified by evaluating $A_{\text{det.}}^{\text{ProbNN}}$ over a range of ProbNN cuts, shown in Fig. 4.21, wherein the ProbNN cuts are parameterised by an arbitrary parameter $\varepsilon \in \{0.1, 0.3, 0.5, 0.7, 0.9\}$ such that cuts of $\text{ProbNN}\pi < \varepsilon$ and $\text{ProbNN}k > 1 - \varepsilon$ are applied to K^\pm , and $\text{ProbNN}k < \varepsilon$ and $\text{ProbNN}\pi > 1 - \varepsilon$ are applied to π^\pm . As the loose and strict ProbNN requirements applied in this analysis correspond to $\varepsilon = 0.4$ and $\varepsilon = 0.2$, respectively, a systematic uncertainty is assigned as half of the difference between $A_{\text{det.}}^{\text{ProbNN}}$ for the $\varepsilon = 0.1$ and $\varepsilon = 0.5$ cases, *i.e.*, the values of ε which bracket those used in the nominal selection.

For $A_{\text{det.}}^{\text{L0}}$, the choice of reference tables is considered as a source of systematic uncertainties. To quantify this, the uncertainty due to the statistics of the reference samples is assigned as a systematic uncertainty on $A_{\text{det.}}^{\text{L0}}$.

Finally, four systematic effects are considered for $A_{\text{det.}}^{\text{track}}$: the choice of weighting approach, the choice of signal model in fits to the calibration samples, the hadronic interactions of pions and the beam-crossing angle.

A systematic due to the choice of weighting approach applied for $A_{\text{det.}}^{\text{track}}(K^- \pi^+)$ is assigned by computing $A_{\text{det.}}^{\text{track}}$ using `BinsReweighter` and taking half the difference against the value of $A_{\text{det.}}^{\text{track}}$ computed using `GBReweighter`.

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays




Figure 4.21: Dependence of the $A_{\text{det.}}^{\text{ProbNN}}$ on the choice of requirements **ProbNNk** and **ProbNNpi**, parameterised by ϵ . Values of ϵ corresponding to the PID requirements applied within and outside the resonant D_s^- phase-space are marked by hollow and solid circles, respectively. The uncertainties on $A_{\text{det.}}^{\text{ProbNN}}$ are solely statistical, and are entirely correlated with one another as they study $A_{\text{det.}}^{\text{ProbNN}}$ within the same sample.

The effect of the choice of signal model is quantified by computing $A_{\text{det.}}^{\text{track}}$ replacing the DSCB model with the sum of two Gaussians and again taking half the difference to the value of $A_{\text{det.}}^{\text{track}}$ obtained using a DSCB.

The crossing angle of the proton beams may induce a phase-space asymmetry with a dependence on the magnet polarity, resulting in a contribution to the detection asymmetry estimated at -0.2% for pions with $p > 20 \text{ GeV}/c$ and $\eta > 4$, and considered significantly smaller for pions outside these ranges and for particles arising from secondary decays, *i.e.*, from the D_s^- [120]. In $B_s^0 \rightarrow D_s^- \pi^+$ combined across each year and magnet polarity, $(11.220 \pm 0.023)\%$ and $(12.34 \pm 0.09)\%$ of associated pions have $p > 20 \text{ GeV}/c$ and $\eta > 4$, in the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively. Taking the product of the detection asymmetry at high p_T and η , and fraction of events for which this applies, this effect is therefore estimated at 2×10^{-4} , which is assigned as a systematic on $A_{\text{det.}}^{\text{track}}$.

As the single-pion tracking efficiencies are computed from muons they neglect any hadronic interactions which pions would be subject to. A systematic uncertainty of 0.04% is assigned to resulting single-pion tracking asymmetries in previous analyses [31]. In the $KK\pi\pi$ final state, the hadronic interaction of the π^+ from the D_s^+ is accounted for in **KPiAsymmetryTool**, so this effect is only relevant for the associated pion, and the systematic uncertainty of 0.04% is thus assigned to $A_{\text{det.}}^{\text{track}}(KK\pi\pi)$. In the $\pi\pi\pi\pi$ final state, this effect is assumed to cancel, such that no corresponding systematic uncertainty need assigned.

The systematic uncertainties assigned to each component of $A_{\text{det.}}$ and their combination as a quadrature sum are given in Table 4.15. The largest contribution in both final states arises from the difference in the asymmetry induced by the kinematic selections when incorporating PID and L0 efficiencies. Additionally, in the $KK\pi\pi$ final state, the systematic uncertainty assigned to account for asymmetries in the hadronic interactions of pions provides a large contribution. Whilst the systematic uncertainties assigned to $A_{\text{det.}}$ are larger than those assigned to A_{raw} , they remain approximately a factor of 10 below the statistical sensitivity to A_{raw} in each final state.

4.7 Consistency of $\langle A_{\text{untagged}}^s \rangle$ in data subsamples

Table 4.15: Summary of systematic uncertainties on $A_{\text{det.}}$.

Component	Systematic	Final state	
		$KK\pi\pi$	$\pi\pi\pi\pi$
$A_{\text{det.}}^{\text{sel.}}$	Per-event weights	4×10^{-4}	2×10^{-3}
$A_{\text{det.}}^{\text{ProbNN}}$	Binning scheme	2×10^{-4}	2×10^{-4}
	ProbNN cuts	3×10^{-4}	5×10^{-4}
$A_{\text{det.}}^{L0}$	Efficiency tables	3×10^{-5}	1×10^{-4}
$A_{\text{det.}}^{\text{track}}$	Weighting approach	2×10^{-4}	—
	Signal model	2×10^{-4}	2×10^{-4}
	π^\pm hadronic interactions	4×10^{-4}	—
	Beam-crossing angle	2×10^{-4}	2×10^{-4}
$A_{\text{det.}}$	Total	7.5×10^{-4}	2.2×10^{-3}

4.7 Consistency of $\langle A_{\text{untagged}}^s \rangle$ in data subsamples

As a validation of the determination of $\langle A_{\text{untagged}}^s \rangle$ described above, the measurement is performed in subsamples of the dataset and compared to the nominal values obtained in the full samples. In particular, subsamples are constructed from regions of the D_s^- phase space, and bins of the B_s^0 and associated pion kinematics. Subsamples from the D_s^- phase space regions are taken as the candidates lying within each individual resonance window (per Section 4.1.1), and the non-resonant region. Subsamples from the B_s^0 kinematics are defined by dividing candidates into four bins depending on whether the corresponding $p_T(B_s^0)$ and $\eta(B_s^0)$ are above or below boundaries $p_T = 8 \text{ GeV}/c$ and $\eta = 3$. A similar split is performed for the associated pion subsamples, using the same η boundary and taking the p_T boundary to be $p_T = 4 \text{ GeV}/c$.

Measurements of $\langle A_{\text{untagged}}^s \rangle$ are obtained in each subsample, following the same procedures as the nominal measurement, for each data-taking year and magnet polarity. The differences between the luminosity-averaged values in each subsample and the nominal value of $\langle A_{\text{untagged}}^s \rangle$ determined in the 2016-18 magnet polarity-combined samples are reported in Table 4.16, where the uncertainties stated are entirely statistical.

The values of $\langle A_{\text{untagged}}^s \rangle$ measured in each subsample are generally consistent with the nominal values, with only a few subsamples in which the listed differences are inconsistent with 0. Where tensions do arise, these are small, with the largest, a difference of -0.135 ± 0.082 , arising in the low $p_T(\pi_{B_s^0}^+)$, high $\eta(\pi_{B_s^0}^+)$ subsample in the $\pi\pi\pi\pi$ final state. Additionally, these differences cancel between the subsamples when combined according to their relative abundance in the full samples, indicating closure of $\langle A_{\text{untagged}}^s \rangle$. It is concluded that the determination of $\langle A_{\text{untagged}}^s \rangle$ (and its constituents) is robust.

4.8 Summary of results and conclusion

Combining the raw and detection asymmetries determined above, the values of $\langle A_{\text{untagged}}^s \rangle$ determined in each final state are:

$$\langle A_{\text{untagged}}^s \rangle_{KK\pi\pi} = (-1.2 \pm 6.2 \pm 1.0) \times 10^{-3},$$

$$\langle A_{\text{untagged}}^s \rangle_{\pi\pi\pi\pi} = (-0.4 \pm 1.6 \pm 0.2) \times 10^{-2},$$

4 Measuring the time-integrated untagged CP asymmetry of $B_s^0 \rightarrow D_s^- \pi^+$ decays

Table 4.16: Differences between the nominal values of $\langle A_{\text{untagged}}^s \rangle$ and values determined in subsamples of the datasets.

Category	Subsample	$\langle A_{\text{untagged}}^s \rangle$ difference	
		$KK\pi\pi$ final state	$\pi\pi\pi\pi$ final state
D_s^- phase space	$\phi \rightarrow K^+ K^-$	0.006 ± 0.019	—
	$K^{*0} \rightarrow K^- \pi^+$	-0.011 ± 0.019	—
	$f_0(980) \rightarrow \pi^+ \pi^-$	—	-0.097 ± 0.044
	Non-resonant	0.003 ± 0.031	0.001 ± 0.046
B_s^0 kinematics	Low p_T , low η	0.028 ± 0.028	-0.039 ± 0.079
	Low p_T , high η	-0.005 ± 0.024	-0.063 ± 0.057
	High p_T , low η	-0.015 ± 0.018	-0.017 ± 0.049
	High p_T , high η	-0.001 ± 0.049	-0.09 ± 0.11
$\pi_{B_s^0}^+$ kinematics	Low p_T , low η	0.020 ± 0.025	-0.072 ± 0.067
	Low p_T , high η	-0.039 ± 0.031	-0.135 ± 0.082
	High p_T , low η	0.001 ± 0.017	0.011 ± 0.049
	High p_T , high η	0.004 ± 0.038	-0.036 ± 0.084

where the constituent terms of each value are listed in Table 4.17. No evidence of NP is observed, with the measured values deviating from the SM case by 0.19σ and 0.25σ in the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states respectively.

Table 4.17: Summary of the components of $\langle A_{\text{untagged}}^s \rangle$, for the combination of all years and magnet polarities.

Source	Asymmetry / %	
	$KK\pi\pi$ final state	$\pi\pi\pi\pi$ final state
$A_{\text{det.}}^{\text{sel.}}$	$0.031 \pm 0.063 \pm 0.038$	$0.03 \pm 0.24 \pm 0.21$
$A_{\text{det.}}^{\text{ProbNN}}$	$0.021 \pm 0.088 \pm 0.057$	$0.08 \pm 0.34 \pm 0.17$
$A_{\text{det.}}^{L0}$	$-0.169 \pm 0.056 \pm 0.007$	$0.00 \pm 0.21 \pm 0.02$
$A_{\text{det.}}^{\text{track}}$	$0.00 \pm 0.06 \pm 0.05$	$-0.01 \pm 0.02 \pm 0.03$
$A_{\text{det.}}$	$-0.12 \pm 0.11 \pm 0.09$	$0.00 \pm 0.21 \pm 0.22$
A_{raw}	$-0.25 \pm 0.62 \pm 0.05$	$-0.3 \pm 1.6 \pm 0.00$
$A_{\text{prod.}}$	—	—
$\langle A_{\text{untagged}}^s \rangle$	$-0.12 \pm 0.62 \pm 0.10$	$-0.4 \pm 1.6 \pm 0.2$

The measurement is statistically limited, with the statistics effectively reduced twice by L0Hadron trigger selection: first by the L0Hadron trigger efficiencies of $\sim 40\%$ from Section 4.2.3, and again by the HCAL separation cut which is applied solely to enable the L0Hadron trigger efficiencies to be determined and has an efficiency of $\sim 25\%$. If the L0 hardware trigger were to be removed, an improvement the statistical sensitivity of $\sim 3\times$ could be expected. Precisely such improvements in the sensitivity of statistically-limited measurements motivate the removal of the L0 hardware trigger for Run 3 data-taking, a focal point of the next chapter.

5 A 30 MHz LHCb software trigger for Run 3

As demonstrated in the previous chapter, the L0Hadron trigger has proven to be a limiting factor in the sensitivity of many measurements. As will be shown in this section, the L0Hadron trigger must be removed to enable LHCb to capitalise on increases in the delivered instantaneous luminosity from Run 2 to Run 3. The trigger system was therefore redeveloped, centred on a real-time analysis-based approach, wherein HLT1 receives data at the full 30 MHz of the pp collisions.

This chapter discusses the upgrade of the LHCb trigger for Run 3, and the upgrade of the LHCb detector as a whole. A data-driven method to evaluate trigger efficiencies is then laid out, along with an accompanying common software tool, `TriggerCalib`. Finally, this method is applied to determine the efficiencies of inclusive HLT1/HLT2 selections.

5.1 The LHCb experiment during Run 3 of the LHC

The LHCb detector and trigger underwent significant upgrades during Long Shutdown 1 (LS1), *i.e.*, between Runs 2 and 3, to facilitate higher energy and higher intensity LHC collisions in Runs 3 and 4 [121]. The centre-of-mass energy of the pp collisions is increased from 13 TeV to 13.6 TeV and the instantaneous luminosity in LHCb is increased from $\mathcal{L}_{\text{inst}} = 4 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$ to $\mathcal{L}_{\text{inst}} = 2 \times 10^{33} \text{ cm}^{-2} \text{ s}^{-1}$ [122]. This increase in $\mathcal{L}_{\text{inst}}$ raises the average number of visible pp interactions per bunch crossing from $\mu_{\text{vis}} \sim 1$ in Run 1/2 to $\mu_{\text{vis}} \sim 5$. The detector must therefore be able to record and efficiently process the busier, *i.e.*, higher occupancy, events of Run 3. However, as is demonstrated in Section 5.1, the signal yields of the L0 trigger saturate when increasing the instantaneous luminosity beyond that of Run 1. The upgrade therefore necessitates the removal of the L0 trigger, with the HLT being upgraded to enable high-quality reconstruction and particle identification in real-time, *i.e.*, the 30 MHz visible interaction rate.

5.1.1 Upgrades to the LHCb detector

The LHCb detector was significantly modified to optimise performance under the busier Run 3 data-taking conditions. The upgraded detector, shown in Fig. 5.1, features a full replacement of tracking systems and upgrades to RICH1 and RICH2, whilst also seeing the SPD, the PS and M1 removed. These upgrades enable the detector to disentangle the multiple visible interactions and efficiently reconstruct the consequent events.

All three trackers of the Run 1/2 LHCb detector are replaced: the VELO by the VELO upgrade [123], T1-T3 by the Scintillating Fibre (SciFi) Tracker, and the TT by the Upstream Tracker (UT) [61].

The semi-circular VELO sensors are replaced with hybrid pixel detector modules [123]. The 52 L-shaped modules, each consisting of $55 \mu\text{m} \times 55 \mu\text{m}$ square pixel sensors, are installed either side of the beam and arranged to form a square aperture 5.1 mm around the beam when the VELO is closed. The resulting resolutions of PV positions in the x and z directions are shown in Fig. 5.2 in comparison to those of Run 2 [124]. PVs can be resolved more precisely in both directions by the upgraded VELO, even for collisions with many more PVs than in Run 2.

5 A 30 MHz LHCb software trigger for Run 3




Figure 5.1: Side-on schematic of the LHCb detector in Run 3 of the LHC from Ref. [122].

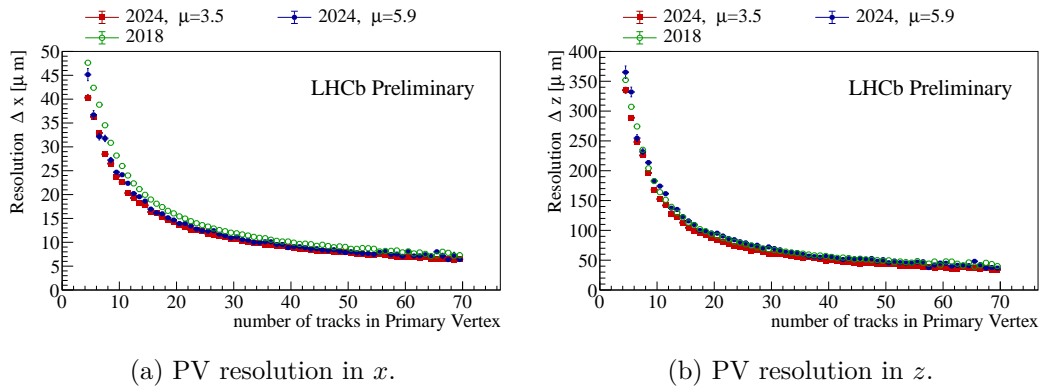


Figure 5.2: Resolutions of reconstructed PV positions in 2018 and 2024, taken from Ref. [124].

The UT replaces the TT, consisting of four planes of vertical staves, with silicon strip sensors on both back and front [61], in the same geometry as the TT. As for the TT, the UT differs between the inner and outer regions, employing narrower and shorter strips in the innermost regions of each layer to achieve a higher granularity.

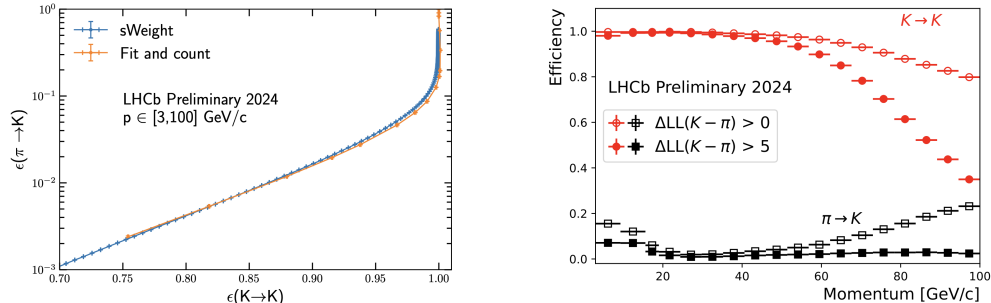
The SciFi tracker consists of three stations, replacing T1-T3, each of which is formed from four layers of mats containing 250 μm -thick scintillating plastic fibres. Light produced by the scintillating fibres is detected by silicon photomultipliers (SiPMs) at the end of the fibres. Both the UT and SciFi maintain the X-U-V-X geometries of their predecessors.

The dipole magnet, which was not replaced, enables a momentum resolution compatible with that of Run 2 (shown in Fig. 3.4a in Section 3.2.1), as reported for 2024 data-taking in Ref. [125].

Less radical changes are made to the PID and calorimetry systems, though upgrades are nonetheless required to enable 30 MHz readout. The HPDs of RICH1 and RICH2 are replaced with multianode photomultiplier tubes and the mirrors of RICH1 are replaced with those of a longer focal length [126]. The SPD, PS and M1, which all had primary purposes of aiding the L0 trigger, are made obsolete by the removal of L0. These subdetectors are thus removed, with the additional benefit of reducing the detector

material budget. The ECAL, the HCAL and M2-M5 require only changes to detector readout, and are left otherwise undisturbed.

The performance of the PID in 2024 data is demonstrated in Ref. [127], from which the efficiency of PID_k is shown as a function of fake-rate and momentum in Fig. 5.3.



(a) Efficiency against fake-rate (as π) of K^\pm identification. (b) Momentum dependence of K^\pm identification efficiency and fake rate.

Figure 5.3: Performance of K^\pm identification in early 2024 data taking, from Ref. [127].

In addition to the upgrades to tracking, particle identification and calorimetry systems, the sensors of the BCM are replaced and the support structure rebuilt. The Probe for LUminosity MEasurement (PLUME) luminometer is also installed, using photomultiplier tubes (PMT) to measure the Cherenkov light of charged particles to estimate the average μ_{vis} and by extension, the average luminosity per bunch crossing [128].

5.1.2 Upgrade of the LHCb trigger

At the heart of the LHCb upgrade is the transition of the trigger to an entirely software-based, real-time analysis approach. To realise this first-of-its-kind approach at LHCb, the L0 trigger is removed and HLT1 is run on graphics processing units (GPUs) to achieve the necessary throughput of 30 MHz.

The strict requirements on p_T and E_T made in L0 to achieve a $40\times$ reduction in the event rate resulted in low efficiencies for the selection of hadrons, electrons and photons, as shown previously in Fig. 3.11. Hadronic decays suffer particularly from these requirements, as is demonstrated throughout Chapter 4, specifically in Section 4.2.3. At higher instantaneous luminosities, the higher occupancy of the SPD and PS further degrades the ability to trigger efficiently on hadrons, electrons and photons. This effect results in diminishing returns on the yields of such decays of interest with increasing instantaneous luminosity, as shown in Fig. 5.4.

The upgraded trigger design is shown in Fig. 5.5, with the flow of data through the trigger shown in Fig. 5.6.

The EB, HLT1 and Allen

The front-end readout of each subdetector, approximately 4 TB/s is collected by the event builder (EB) CPU farm, consisting of 163 nodes. These nodes contain TELL40 FPGA readout boards which are responsible for the pre-processing of raw information from the subdetectors into multi-fragment packets for the EB to group by run bunch crossing into multi-event packets (MEPs) of ~ 1000 events [131]. These MEPs are passed to the EFF in which HLT1 runs, which in Run 3 consists of GPUs hosted in-situ on the




Figure 5.4: Signal yields of the LHCb trigger from Run 2 at instantaneous luminosity working points, normalised to the yields at the Run 1 instantaneous luminosity of $\mathcal{L}_{\text{inst}} = 2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$. from Ref. [122].

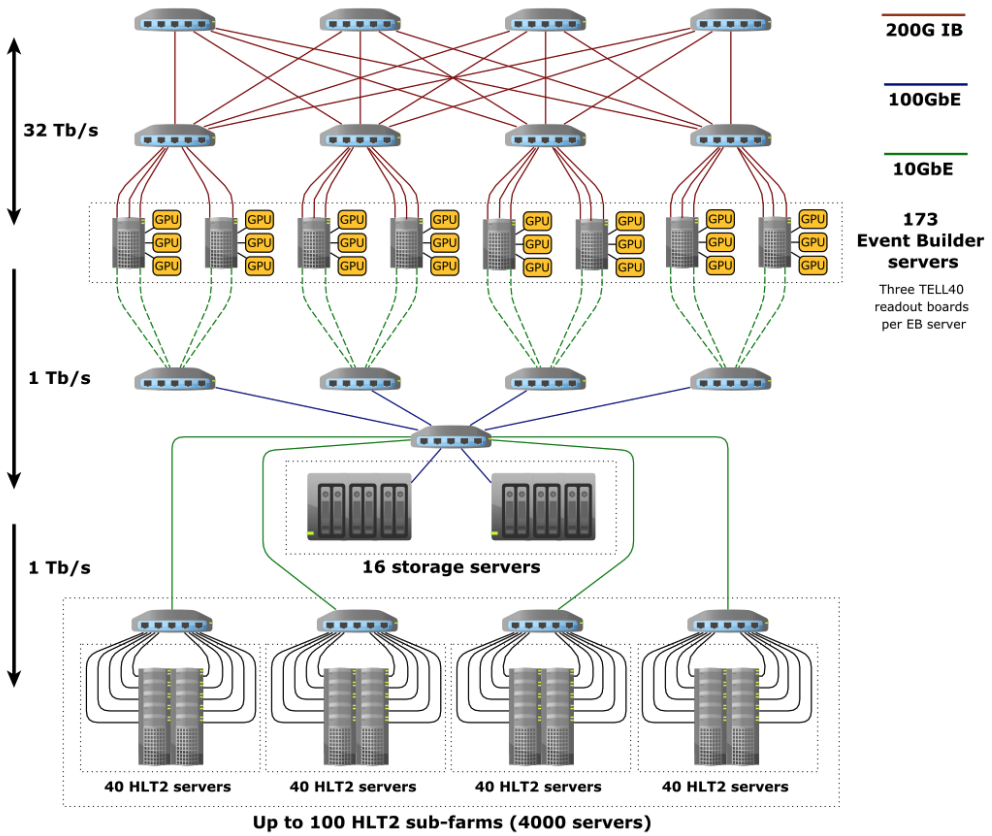


Figure 5.5: Schematic diagram of the Run 3 LHCb trigger configuration from Ref. [129].

PCIe slots of the EB nodes, removing the need to transfer the data outside of the EB nodes.

The GPUs of HLT1 perform real-time partial reconstruction (track and vertex reconstruction) and selection. At the core of the redesign is the need for HLT1 to be

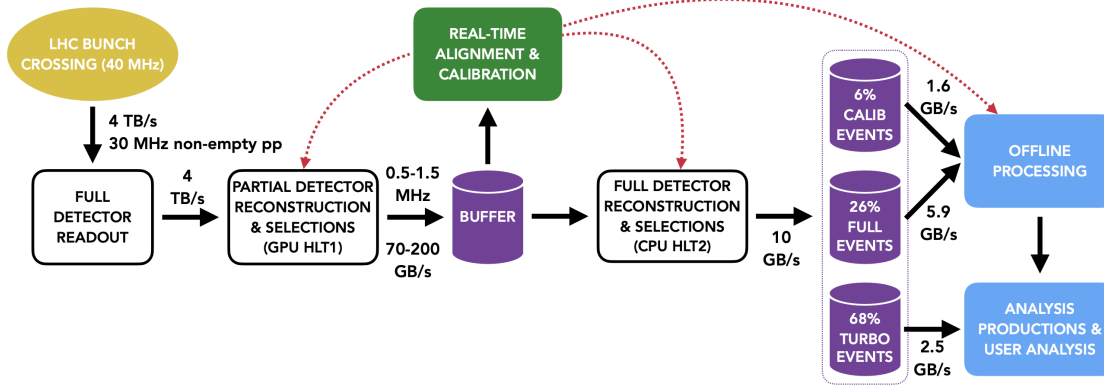


Figure 5.6: Schematic diagram of the flow of data through the LHCb trigger in Run 3 from Ref. [130]. The onward flow of data in offline processing is shown in more detail in Fig. 5.10.

architecture-agnostic, (and, by extension, supplier-agnostic), and built around off-the-shelf consumer GPUs. The former ensures a robustness against the ever-turbulent GPU market [132], whilst the latter ensures that the trigger can be scaled up in a cost-effective manner. Extensive comparisons of potential GPUs were carried out ahead of the upgrade, *e.g.*, those shown in Fig. 5.7, resulting in the adoption of NVIDIA RTX A5000 GPUs [133]. Each EB node has 8 PCIe slots, 5 are used for data input/output, with the remaining 3 available for the GPUs. Of these 3 slots per node, one GPU was already in place prior to the start of Run 3 data-taking in 2022, a second GPU was added to each node ahead of 2023 data taking, and a third in August 2024, bringing the total number of GPUs up to almost 500.




Figure 5.7: Comparison of Allen throughputs for GPUs considered for HLT1 from Ref. [129].

The Allen software framework run on these GPUs was designed to provide a highly-parallelised GPU-compatible implementation of HLT1 which can be operated at scale [133]. The partial reconstruction of objects is an embarrassingly parallel problem [134], in that events can be treated separately, as can their constituent objects. This enables an efficient mapping of track and vertex reconstruction onto GPUs.

Significant effort has gone into developing and optimising the reconstruction algorithms to ensure high performance, whilst operating within the narrow throughput constraints

imposed by design. For example, the Looking Forward algorithm, optimised to run on GPU, reconstructs tracks by defining search windows in each SciFi layer based on extrapolated VELO- and UT-tracks. Triplets of hits are sought within the search windows, which are extended to the other layers to be used in a track fit, the results of which are filtered based on track quality [135]. The efficiency of this reconstruction and subsequent throughput at varying numbers of PVs are shown in Fig. 5.8.

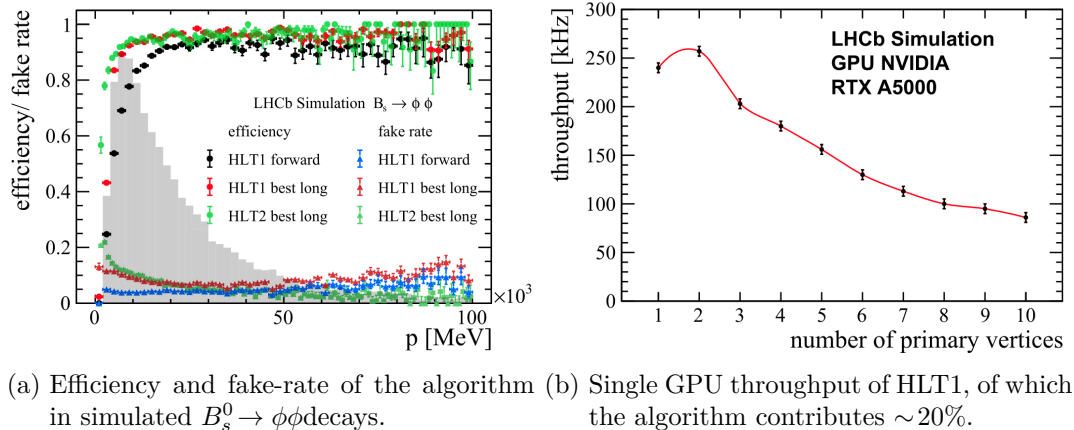


Figure 5.8: Performance of the Looking Forward track reconstruction algorithm on simulated samples, as detailed in Ref. [135].

The high-quality reconstruction of tracks and vertices, and access to calorimetry and muon information, enables powerful selections in HLT1, *i.e.*, highly efficient selections with strong background suppression. As in Run 2, these are typically inclusive selections; however, in Run 3, HLT1 runs a handful of exclusive selection algorithms, *e.g.*, those targeting $D^0 \rightarrow K^- \pi^+$ and $\Lambda \rightarrow p \pi^-$. The output of these trigger lines must be distributed equitably across the LHCb physics interests within the limited output bandwidth. This allocation is determined by optimising thresholds on the HLT1 selection to maximise the efficiency of HLT1 across a range of signal channels, with this procedure discussed in more detail in Ref. [136]. Efficiencies of the HLT1 selection algorithms, namely HLT1TrackMVA and HLT1TwoTrackMVA, are deferred to Section 5.2.3, where these are discussed in detail.

HLT1 and Allen have been extended beyond their design functionality, *e.g.*, in the reconstruction of downstream tracks [137] and the clustering of ECAL hits [138]; however, these topics are beyond the scope of this thesis.

Events selected by HLT1 are saved to a 40 PB disk buffer. Under nominal conditions, events are written to this buffer at 70-200 GB/s, providing a minimum of ~ 2.5 days in which to process these events.

Real-time alignment, calibration and monitoring

As in Run 2, the events saved to the disk buffer are used to perform the alignment and calibration of sub-detectors in real time. This is performed on a similar timescale to that of Run 2; however, improvements to the reconstruction throughput reduce the computational needs of the alignment and calibration, allowing these resources to instead be used for HLT2 processing and for monitoring [139].

With the higher luminosity and subsequent event rate in Run 3, the use of monitoring to provide performance insights and identify issues as in data-taking as quickly as possible takes on a greater importance. The move to a fully software-based trigger enables the monitoring of high-level performance quantities, such as tracking and particle identification efficiencies, on a per-run/fill basis. These quantities provide a more detailed insight into the detector and trigger performance, as these are analysis-level quantities with the real correlations to their dependent properties. These quantities also reflect the performance of the first stages of the data flow, and thus are necessary to identify and act upon issues in a timely manner.

The monitoring of high-level performance quantities is implemented in the **CalibMon** monitoring framework, to which the author contributed significantly to the development, configuration and maintenance. This framework, first implemented in 2024, has been used throughout Run 3 to identify and understand issues in data-taking in real time. The framework, the quantities studied and its results from early 2024 data-taking are discussed below.

The **CalibMon** framework takes events from the HLT1 output buffer in real time and applies a set of dedicated HLT2 lines in a HLT2-like processing to fill monitoring histograms on which further analysis, *e.g.*, fitting, can be performed and the results displayed. **CalibMon** is configured to run with 20 threads per node across 10 HLT2 nodes, receiving events from HLT1 at $\mathcal{O}(1\text{ kHz})$.

This framework allows for the study of tracking and PID efficiencies using dedicated HLT2 Turbo lines to select decay channels of relevance to each method.

An example of a **CalibMon** monitoring page, that of the SciFi HLT2 tracking efficiencies from $K_S^0 \rightarrow \pi^+\pi^-$, in the monitoring web-application **Monet** [140] is shown in Fig. 5.9.



Figure 5.9: Monitoring web page for the **CalibMon** monitoring of $K_S^0 \rightarrow \pi^+\pi^-$ for a single LHC fill in early 2024 data-taking.

The HLT2 tracking efficiencies in both $J/\psi \rightarrow \mu^+\mu^-$ and $K_S^0 \rightarrow \pi^+\pi^-$ decays are estimated using the tag-and-probe approach discussed in Section 4.2.4, wherein one μ^\pm/π^\pm

¹Unlike in Section 4.2.4, the J/ψ candidates are not required to originate from a b -hadron decay.

is reconstructed as a long track, and the other as a partial track which is then matched against the long tracks of the event. In $J/\psi \rightarrow \mu^+\mu^-$, as in Section 4.2.4, these partial tracks are formed from hits in the VELO or SciFi and in the muon station, for evaluating the tracking efficiencies with respect to the SciFi and VELO, respectively. If the UT hits are available, these are also used in both types of partial track. Two $\mu^+\mu^-$ mass histograms are filled per probe track type and charge, one before and one after probe track matching is performed. A simultaneous fit is performed across each pair of mass histograms, with a DSCB (Eq. A.1 in Appendix A.3) modelling signal and an exponential function models combinatorial background, with all parameters shared between the two histograms except for the yields of the components, from which the efficiency is computed. This approach is known as fit-and-count. As only $\mathcal{O}(10^3)$ signal candidates can be expected to be collected per fill, the efficiencies are computed globally, *i.e.*, with no kinematic binning applied.

K_S^0 are produced much more abundantly than J/ψ , so by also studying the decays of $K_S^0 \rightarrow \pi^+\pi^-$, the additional statistics can be used to compute tracking efficiencies in kinematic bins and subsequently also the highly kinematic-dependent tracking asymmetries per Eq. 4.24. As these decays do not involve the muon stations, it is only possible to study the SciFi tracking efficiency, taking the probe tracks to be VELO tracks. For each probe track charge, and their combination, mass histograms are filled in specific ranges of the probe π^\pm momentum, transverse momentum and pseudorapidity, such that the signal yield is evaluated in kinematic bins. Whilst the samples of $K_S^0 \rightarrow \pi^+\pi^-$ decays collected are already very pure, a sideband-subtraction approach, as per Eqs. 5.20 & 5.21 in Section 5.2.2, is applied to extract the signal yields.

Efficiencies of μ^\pm and h^\pm PID requirements are evaluated using the same tag-and-probe methods as described in Section 4.2.2, allowing the effect of any issues in the muon stations to be identified in quasi-real time. For the identification of muons, $B^+ \rightarrow J/\psi (\mu^+\mu^-) K^+$ samples are collected, where `isMuon` (defined in Section 3.2.2) is required for the tag muon and not required for the probe muon. The efficiency is then taken as the fraction of probe muons passing an `isMuon` requirement, with the signal yields obtained by a fit-and-count approach applied to the dimuon mass using the same model as for the $J/\psi \rightarrow \mu^+\mu^-$ tracking efficiencies. Sufficiently many $B^+ \rightarrow J/\psi (\mu^+\mu^-) K^+$ candidates are anticipated that, similarly to the $K_S^0 \rightarrow \pi^+\pi^-$ tracking efficiencies, the muon identification efficiencies are computed in bins of probe muon p . The efficiencies are evaluated both separately for probe muons with $p_T < 800$ MeV and $p_T > 800$ MeV, and global efficiency, *i.e.*, making no p or p_T distinctions.

Whilst the effect of changes in the muon stations on the muon identification efficiencies can usually be understood through monitoring of other quantities such as the muon hit efficiencies, the same is not true for the efficiency of PID requirements on hadrons, which are non-trivially related to the performance of RICH1 and RICH2. It is therefore vital to include the hadronic PID efficiencies in `CalibMon`, as this provides a unique insight into whether the PID systems are working together as intended. The efficiencies of requirements on PID_k on K^\pm and π^\pm , PID_p on p and π^\pm , and PID_{pk} on p and π^\pm are determined from $D^{*+} \rightarrow D^0\pi^+$ and $\Lambda \rightarrow p\pi^-$ decays, in the same approach as for `PIDCalib2`, described in Section 4.2.2 [104]. These efficiencies are evaluated at working points of $\text{PID}_i > x$ of $x \in \{-10, -5, 0, 5\}$. A similar method is also applied to the electron PID efficiencies, in which an analogous tag-and-probe approach uses the dielectron system of $B^+ \rightarrow J/\psi (e^+e^-) K^+$ decays to evaluate the efficiency of $\text{PID}_e > x$ at working points of $x \in \{0, 5\}$.

HLT2 in Run 3

The upgraded HLT2 remains largely similar in function to the Run 2 HLT2, performing full offline-quality event reconstruction and subsequent event selection. HLT2 is still run on a CPU farm, though operating with significantly higher input and output rates, of 1 MHz and 100 kHz, respectively.

Events are divided into streams depending on their event model and purpose: events with full persistency are saved to the Full stream, events using the Turbo model [87] are saved to the Turbo stream, events intended for offline calibration (which are Turbo events with specific sub-detector raw banks attached) are saved to the Calibration (TurboCal) stream. A dedicated NoBias stream saves a small number of events without imposing a trigger requirement, as these are used for detector studies. Data sent to the Turbo/TurboCal streams, which are typically accessed regularly and may be needed at short notice, are saved to hard drive disks. Data sent to the Full stream, which are typically only accessed for further processing, are saved to magnetic tape [141].

As of 2025, HLT2 contains over 3000 trigger lines. The use of the Turbo event model is expanded compared to Run 2, to cover many exclusive selections. Many of these selections attach additional information such as calorimeter raw banks or additional tracks for isolation studies. Most inclusive selections remain fully persisted, with further processing applied offline to reduce candidates to Turbo-like events.

As for HLT1, discussion of efficiencies of inclusive HLT2 selections are deferred to Section 5.2.3. Efficiencies of exclusive selections are, by construction, analysis-specific, though these are discussed briefly in the context of $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^+$ decays in Section 6.1.1.

Offline data processing and simulation

The flow of data through the offline processing, which leads on directly from the online data flow shown in Fig. 5.6, is shown in Fig. 5.10. The primary aim of this offline processing is to reduce the footprint of the data recorded by the trigger and to provide these data to analysts [142].

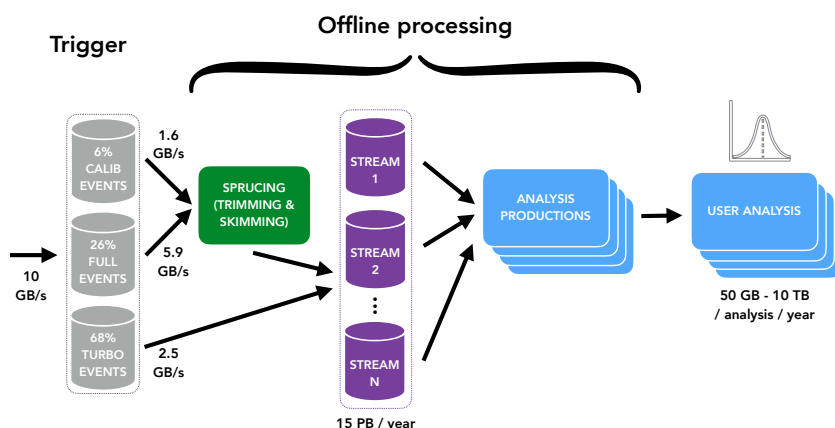


Figure 5.10: Dataflow of the LHCb offline processing from Ref. [130].

The events recorded by the trigger undergo further selection (skimming) and reduction in event content (slimming), known as Sprucing, saving Turbo-like events to disk storage [143]. The output of Sprucing and the HLT2 Turbo stream are saved to disk as a

series of streams based on the working group and/or physics content. These events can be further processed and written into *Root* NTuples for offline user analysis through the Analysis Productions framework, which orchestrates the distribution of this processing for analysts [144, 145].

Making use of improvements to the underlying *Gaudi* framework necessitated by the trigger upgrade, the simulation software was consolidated into the experiment-independent *Gaussino* framework [146]. This framework provides a mechanism to allow integration with event generators beyond *Pythia* and *EvtGen*, and supports descriptions of detector geometries through *DD4HEP* [147].

5.1.3 Status of Run 3 data-taking

In the course of Run 3 thus far, LHCb has collected over 17 fb^{-1} of pp collisions, as shown per-year in Fig. 5.11. This data-taking has seen the detector and trigger upgrades commissioned, operating at, and in many cases beyond their design capabilities. This significant increase in the total integrated luminosity collected by LHCb provides vast statistics for physics analyses, the first of which, a measurement of c -meson production asymmetries in 2022 and 2023 data, has already been submitted for publication [148]. However, the data-taking has not been without incident; in January 2023, vacuum protection equipment failures caused damage to the aluminium foil isolating the VELO from the LHC vacuum [149]. Whilst the VELO modules were undamaged, the VELO could not be fully closed for the duration of the 2023 data-taking.




Figure 5.11: Cumulative integrated luminosity of pp collisions recorded by LHCb, shown by year of data-taking [150].

The remainder of this thesis focuses on the pp collision data recorded in 2024 and 2025, up to the technical stop in June 2025. Blocks of data are labelled within these samples, each of which contains data taken with the same magnet polarity and trigger configuration, *e.g.*, a stable HLT1 and alignment configuration. The magnet polarities and integrated luminosities of the data blocks considered in this thesis are listed in Table B.1 in Appendix B. Data taken in 2022 is not discussed as the configuration of the trigger has since changed significantly. Data taken in 2023 is not discussed (with the exception of

monitoring histograms shown in Section 5.2.1) as the detector conditions do not reflect those of nominal data-taking.

5.2 Development and performance of selections for the Run 3 software trigger

The upgrade of the LHCb trigger necessitates large-scale redevelopment of the selection algorithms employed in HLT1 and HLT2. In HLT1 these selections were developed from scratch along with the `Allen` framework, with these developments focusing on efficiently performing these selections on GPU architectures [133]. Given the large increase in the HLT2 input bandwidth, the selections in `Moore` required redevelopment, leveraging the offline-quality reconstruction enabled by the upgrade and making broader use of the Turbo event model and data stream [62]. The author of this thesis made significant contributions to developments in HLT2, as discussed here.

This section provides a summary of the development process for HLT2 trigger lines, illustrated by the examples of the cut-based inclusive detached dilepton trigger development and topological b -hadron trigger retuning. The data-driven evaluation of trigger efficiencies is then discussed in detail, with these methods then applied to determine trigger efficiencies of HLT1 and HLT2 in 2024 and 2025 data-taking.

5.2.1 Development of inclusive triggers for HLT2

With the Run 3 HLT2 receiving events from HLT1 at a rate $\sim 10\times$ greater than in Run 2, the selections in HLT2 were retuned and redeveloped. In general, exclusive HLT2 and Sprucing lines apply requirements close to those applied in offline physics analyses, often inheriting many of their selection requirements from corresponding stripping lines in Run 2. As exclusive lines target one or a handful of specific decays, their output rate can be tuned to a value appropriate for the branching fraction(s) of the targeted decay(s). Inclusive triggers, on the other hand, must strike a balance between sufficiently strict requirements to suppress background(s) and sufficiently loose enough requirements to provide high efficiencies across a broad range of signals. The inclusive triggers also provide a safety net to ensure that interesting events are selected even if they do not have a corresponding exclusive line, *e.g.*, if a physical analysis is devised after the start of data-taking. Nonetheless, the procedure by which these trigger lines are developed is common between exclusive and inclusive selections.

Trigger lines in `Moore` are specified in Python, with this interface designed such that analysts with a limited software development background can contribute, applying their expertise in the selections of their corresponding analyses. Lines are developed according to five key quantities which must be studied thoroughly: signal efficiency, signal purity, output rate, output bandwidth, and throughput. The signal efficiency and purity of a trigger line are closely related: the efficiency is defined as the fraction of visible signal events, *i.e.*, those with reconstructible final state particles within the detector acceptance, selected by the line; the purity is defined as the proportion of all candidates selected by the line which correspond to signal events. The output rate and bandwidth are even more directly related: the output rate is the number of events selected per second; the output bandwidth is the amount of output data produced per second, which, if each selected event is of a similar size, is simply the product of the output rate and average event size. Finally, the throughput is defined as the number of input events which can be processed by the line per second. Requirements on combinations of these

quantities (dependent on the physics case of the lines) must be met before lines are deployed in HLT2.

The signal efficiency can be determined by studying the fraction of MC simulated signal events which are selected by the trigger:

$$\varepsilon^i = \frac{N_{\text{out}}^i}{N_{\text{in}}^i}, \quad (5.1)$$

wherein $N_{\text{in(out)}}^i$ are the numbers of events in sample i entering (passing) the selection of interest. This treatment is generally sufficient for development of trigger lines, but as discussed in Section 4.2.3, is insufficiently accurate for physics analyses. The data-driven approach used instead is discussed in detail in Section 5.2.2.

The output rate and bandwidth can be studied by processing either minimum bias MC simulated samples, *i.e.*, samples generated to reflect the events produced by the pp collisions without biases towards specific signals, or HLT1-filtered pp collision data samples. Taking the input rate of HLT2, R_{in} , as the output rate of HLT1, the HLT2 output rate can be computed as

$$R_{\text{out}} = R_{\text{in}} \cdot \frac{N_{\text{out}}}{N_{\text{in}}}. \quad (5.2)$$

The output bandwidth can be computed from the resulting compressed file size, s_{out} , as

$$B_{\text{out}} = s_{\text{out}} \cdot \frac{R_{\text{in}}}{N_{\text{in}}}. \quad (5.3)$$

The throughput of the line as a whole and its constituent algorithms are reported by **Moore** after a sample has been processed. Additionally, the MC simulated minimum bias samples, in which the true identities of the particles involved are known, can be used to obtain the signal purity. As processing large minimum bias samples is computationally intensive and time-consuming, signal purity is typically not studied in development.

The trigger lines are not run in isolation, instead **Moore** runs with all lines included in a given HLT2 configuration applied concurrently. Events on which multiple trigger lines fire are recorded only once, saving the corresponding candidates. It is therefore often useful to understand the overlap between lines, as the additional bandwidth usage from a single line will be smaller if other lines are employed which select many of the same events. Overlaps between selections A and B can be quantified in two different ways: firstly, in the Jaccard index [151],

$$J(A, B) = \frac{|A \cap B|}{|A \cup B|}, \quad (5.4)$$

and secondly, in the conditional probabilities

$$P(A \cap B) = \frac{|A \cap B|}{|A|}, \quad (5.5a)$$

$$P(A \cap B) = \frac{|A \cup B|}{|B|}. \quad (5.5b)$$

Both metrics provide a different approach to understanding the overlaps between selections and their usefulness depends on the use case. For example, between two exclusive lines of a similar output rate, the Jaccard index gives an indication of their similarity,

whilst between an exclusive and inclusive line the conditional probability of the exclusive line firing given that the inclusive line fires is often more useful.

As the quantities discussed above must be determined frequently, both ad hoc in the development of lines and regularly to provide summaries of these quantities over time. A bandwidth testing framework was developed to streamline this, in which the tests of the bandwidth and throughput of HLT2 and Sprucing can be requested based on specific released versions of the LHCb software or development branches of the software [152]. This runs standardised tests centrally, from which the bandwidth and throughput are computed and reported on dashboards. Additionally, a tool was developed to enable estimations of the performance of HLT1 /HLT2 lines in MC simulated samples, named `HltEfficiencyChecker` [153]. This tool, which is built into `DaVinci`, provides an implementation of the output rate, signal efficiency and selection overlap calculations discussed above. These are all evaluated by running HLT1 /HLT2 as necessary on data, minimum bias MC or signal MC samples and tupling the output in a specific format. Whilst the initial implementation of the tool, capable of rate and efficiency calculations, were developed prior to this work, implementation of selection overlap calculations were developed by the author. This infrastructure allows for the test-driven development of trigger lines, wherein the performance of the lines can be monitored throughout development.

The methods described above were applied extensively in the development of two inclusive triggers deployed in HLT2, the cut-based inclusive detached dilepton trigger and the topological b -hadron trigger. The author was responsible for the development of the former and was an important contributor to the retuning of the latter, both of which are discussed below.

Development of the cut-based inclusive detached dilepton trigger

The efficient selection of b -hadron decays to final states containing pairs of charged leptons is essential to LHCb physics analyses, such those described in Section 2.4.2. These decays can be identified by two key attributes: the distinct b -hadron decay topology (per Section 2.2.4) and the presence of the dilepton system, originating from a secondary vertex detached from the PV. A collection of MVA-based inclusive HLT2 lines, known as the MVA-based inclusive detached dilepton ($MVA\ell\ell$) trigger, was developed prior to the start of this work. The $MVA\ell\ell$ trigger applies monotonic Lipschitz NNs [154] to select events containing 2-, 3- and 4-body candidates, *i.e.*, those containing a detached dilepton signature and up to two additional tracks. Lines are included to select candidates containing both $\ell^+\ell^-$ (opposite-sign, OS) and $\ell^+\ell^+$ (same-sign, SS), as the latter are regularly used to study background distributions and offline selections in analyses, *e.g.*, in Ref. [155]. It is crucial that the response of the NN is well-understood, ensuring that these events are selected on a reasonable physical basis, else the NN may miss signal events or introduce sculpting effects in background distributions [156]. To this end, a cut-based equivalent to the $MVA\ell\ell$ trigger, the cut-based inclusive detached dilepton ($CB\ell\ell$) trigger, was developed.

The $CB\ell\ell$ trigger consists of lines selecting 2-, 3- and 4-body candidates containing an OS or SS dilepton signature, mirroring those of the $MVA\ell\ell$ trigger. As the $CB\ell\ell$ trigger is intended for validation of the $MVA\ell\ell$ trigger (rather than for event selection alone), a stricter cut-based selection can be performed in the $CB\ell\ell$ trigger than is performed as a preselection on $MVA\ell\ell$ trigger candidates prior to applying the NNs. These stricter selection requirements, primarily on the kinematic and topological candidates of the final state particles and their combinations, result in a lower signal efficiency than in the

MVA $\ell\ell$ trigger, but are necessary to ensure an acceptable output rate, *i.e.*, no higher than the MVA $\ell\ell$ trigger. If both the CB $\ell\ell$ and MVA $\ell\ell$ triggers are operating correctly, and the MVA $\ell\ell$ NNs are well-trained, then all of the events selected by the real CB $\ell\ell$ trigger lines should also be selected by the MVA $\ell\ell$ trigger. If this is the case, then the real CB $\ell\ell$ trigger have no impact upon the HLT2 output rate/bandwidth and efficiencies, instead only affecting the throughput.

Additional CB $\ell\ell$ lines extend this trigger beyond validating the MVA $\ell\ell$ trigger; lines are included in which PID requirements on either or both of the leptons are reversed, dubbed “single-fake” and “double-fake” lines, with lines with non-inverted PID requirements then dubbed as “real”. These lines can be used by analysts to understand the impact of the PID selections applied in the CB $\ell\ell$ trigger and, by extension, MVA $\ell\ell$ trigger, as these lines select the events which would otherwise have been removed by the PID requirements. As reversing the PID requirement(s) removes large quantities of signal events and introduces much larger quantities of background events, and since the non-PID selections of the lines must be identical between real and fake lines for meaningful comparisons to be drawn, the only way to control the output rate of the single- and double-fake lines is to apply prescales to these lines, discarding all but a fraction of selected events.

In the first iteration of the CB $\ell\ell$ trigger, candidates were constructed differently to those of the MVA $\ell\ell$ trigger. Whereas in the MVA $\ell\ell$ trigger the candidates were constructed such that the $\ell\ell$ and additional tracks are contained somewhere in the event, the CB $\ell\ell$ trigger was constructed such that each line corresponded to the whole event, *i.e.*, $b \rightarrow \ell\ell (+X)$. The initial set of selections were taken from a combination of the MVA $\ell\ell$ trigger preselection and the many $b \rightarrow \ell\ell (+X)$ cut-based exclusive lines. For the $e\mu$ lines, the selections on each lepton are taken to be exactly those applied to the muons in the $\mu\mu$ lines and electrons in the ee lines.

From this merged set of cut-based selections, the next step was to tune the selection cuts for the most powerful variables, *i.e.*, those for which a small change in the selection cut provides the greatest change in efficiency/suppression. The most powerful set of variables were the PID variables, as these are responsible for suppression of misidentified backgrounds. Second to these were variables of track and vertex quality, namely χ_{IP}^2 of the leptons and additional tracks, and the χ_{vtx}^2 of the dilepton and b vertices. The tuning was thus carried out by performing 2-dimensional scans in the PID variable cut and track χ_{IP}^2 /vertex χ_{vtx}^2 spaces, with the rates of each line computed in minimum bias MC, and the efficiency of the 2-, 3- and 4-body lines computed on signal MC for 2-, 3- and 4-body b decay channels ($B_s^0 \rightarrow \ell^+\ell^-$, $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^+$ and $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^{*0}$), respectively. By scanning in both dimensions at once, correlations between the two selections can be captured. For example, the scan across the requirements on PID_e and χ_{IP}^2 for the electrons of the 2-body real OS dielectron line is shown in Fig. 5.12. Cuts on the PID variables are conventionally set at integer values, so working points in $\text{PID}_{e,\mu}$ in integer steps between 0 (the loosest present in the lines from which the cuts were inferred) and 4 (generally considered to be an extremely tight selection) were considered.

The set of cut-based selections derived from this tuning, with which the CB $\ell\ell$ trigger operated for the duration of 2023 data-taking, are listed in Table B.2 in Appendix B.1. The single-fake lines were prescaled by a factor of 40, whilst the double-fake lines were prescaled by a factor of 2000.

5.2 Development and performance of selections for the Run 3 software trigger




Figure 5.12: Rates and efficiencies in $B_s^0 \rightarrow e^+e^-$ MC of the 2-body dielectron $CBll$ line for selection working points in PID_e (each curve, stricter for lighter curves) and electron χ_{IP}^2 (each point per curve). Chosen working point is shown by dashed lines.

During 2023 data-taking, monitoring histograms from the output of HLT2 were made available for developers and analysts to verify the performance of their HLT2 lines. This enabled a first chance to look at the $CBll$ and, more importantly, to compare its output with that of the $MVAll$. A small selection of the monitoring histograms from the dimuon lines are shown in Fig. 5.13, with each corresponding $MVAll$ histogram overlaid.

Those of most interest were the 2- and 3-body invariant mass histograms², as some high-statistics signals, such as $J/\psi \rightarrow \mu^+\mu^-$, should be visible above backgrounds. In particular, the two-body invariant mass histograms (Fig. 5.13a) both return sensible dimuon spectra, *i.e.*, the ϕ and $\psi(2S)$ dimuon resonances also visible. However, the $CBll$ trigger is significantly less efficient than the $MVAll$ trigger: from simple fits to the J/ψ resonances, the 2-body $CBll$ trigger selects almost 20 times fewer $J/\psi \rightarrow \mu^+\mu^-$ candidates. A similar inefficiency is seen in the 3-body histograms (Fig. 5.13b), with a similar deficit present with respect to the equivalent $MVAll$ line. Examining the histograms of χ_{vtx}^2 and χ_{IP}^2 for the 2-body lines (Figs. 5.13c and 5.13d, respectively), more significant differences between $CBll$ and $MVAll$ outputs are seen. In addition to the deficits seen in the mass distributions, the distributions of χ_{vtx}^2 and χ_{IP}^2 are both truncated. This truncation is the result of the b -hadron candidate cuts (per Table B.2), with the $CBll$ lines selecting candidates in which the dilepton combination has the b -hadron topology, whilst the $MVAll$ lines select candidates in which the dilepton combination is part of the b -hadron topology. This subtle difference results in 2-body $CBll$ lines configured to reject valid signal candidates with > 2 -body final states. As for the $MVAll$, the 2-body lines should be responsible for the majority of the selection, with the 3- and 4-body lines supplementing this. It was therefore necessary to restructure and retune the $CBll$ trigger ahead of 2024 data-taking.

The retuning of the $CBll$ was a two-step process: the structure of the 2-body lines was changed to align with the $MVAll$ trigger, then the cuts were redetermined to account for the radically different set of 2-body candidates selected. As the change in structure allows the $CBll$ trigger to select dilepton candidates from parts of b -hadron decays,

²The 4-body invariant mass histograms were significantly less populated than the 2- and 3-body histograms.




Figure 5.13: Monitoring histograms from HLT2 processing of Runs 270352-270563, comparing the events selected by the MVA_{ll} (light blue) and CB_{ll} (dark blue) triggers.


it has the secondary effect of increasing the overlap between the 2-body lines and the 3-/4-body lines, permitting looser cuts on the latter.

Firstly, the selection cuts applied to the dilepton candidate were adjusted, p and p_T minima to suppress extremely soft candidates, and the $\chi_{\text{VTX}}^2/N_{\text{dof}}$ and χ_{FD}^2 cuts were loosened. Similarly for the b candidates, the χ_{IP}^2 cut was removed and the $\chi_{\text{VTX}}^2/N_{\text{dof}}$ and χ_{FD}^2 cuts loosened. In addition to the stricter kinematic requirements on the dilepton candidates, the p and p_T minima applied to the leptons and additional tracks were also tightened. Furthermore, by the time of the retuning, two powerful requirements had become sufficiently validated to be used in the inclusive trigger lines: IsMuon, requiring that a track has associated hits in the muon systems, and the ECAL acceptance cut in which tracks are required to be long with InBrem and InECAL, upstream with InBrem or downstream with InECAL. These requirements were imposed on muons and electrons, respectively, providing a significant suppression of background, particularly in the case of electrons.


To determine the extent to which the $\chi_{\text{VTX}}^2/N_{\text{dof}}$, χ_{FD}^2 and track χ_{IP}^2 requirements could be loosened, sets of selection cuts in which these cuts were chosen randomly within reasonable ranges, say $[a, b]$. Since tighter values of these thresholds are more likely to provide an optimal configuration than looser values, a linear probability distribution

5.2 Development and performance of selections for the Run 3 software trigger

$f(x) \propto -(x - b)$ ($f(x) \propto (x - a)$) to draw the random maximum (minimum) values, such that $f(b) = 0$ ($f(a) = 0$). HLT2 was run on HLT1 -filtered MC samples for each set of random selection cuts, with the combined CB $\ell\ell$ trigger rate and efficiency across a range of channels then evaluated with `HltEfficiencyChecker`. As for the 2-dimensional scans of the initial tuning procedure, the rates and efficiencies of each of these selection configurations were plotted to determine the requirements which could be tightened further for a smaller reduction in efficiency. An example of this is shown in Fig. 5.14 (prior to including the ECAL acceptance cut) for efficiencies of the dielectron and dimuon lines in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$ and $B^+ \rightarrow J/\psi(e^+e^-)K^+$, respectively.



(a) Dimuon CB $\ell\ell$ trigger rates plotted against $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$ efficiencies.



(b) Dielectron CB $\ell\ell$ trigger rates plotted against $B^+ \rightarrow J/\psi(e^+e^-)K^+$ efficiencies.

Figure 5.14: Rates and $B^+ \rightarrow J/\psi(\ell^+\ell^-)K^+$ efficiencies of the CB $\ell\ell$ trigger for each random set of topological selection cuts.

The cuts on χ_{IP}^2 of leptons cannot be tightened further without significant impact on the signal efficiencies; whilst the cuts on dilepton χ_{FD}^2 provide some small reduction in rate. In the 2-body lines, the additional `isMuon` and ECAL requirements were sufficient to suppress background, whilst the most effective suppression of background for the 3- and 4-body lines was provided by the χ_{IP}^2 requirement on additional tracks, which was tightened to match the $\chi_{\text{IP}}^2 > 36$ requirement imposed on leptons. The final set of cuts determined by the retuning, as deployed in data-taking from 2024 onwards, are summarised in Table B.3 in Appendix B.1. The same prescales were applied for the single- and double-fake lines as in 2023.

To demonstrate the improvement in performance brought by the CB $\ell\ell$ retuning, the total efficiencies of the 2023 and 2024 CB $\ell\ell$ trigger and 2023 MV $\ell\ell$ trigger were evaluated in $B_s^0 \rightarrow \ell^+\ell^-$, $B^+ \rightarrow J/\psi(\ell^+\ell^-)K^+$, $B_s^0 \rightarrow J/\psi\phi$, $\Lambda_b^0 \rightarrow J/\psi(\mu^+\mu^-)pK^-$ and $\tau^+ \rightarrow \mu^+\mu^-\mu^+$ MC samples, as shown in Fig. 5.15. In every channel, the efficiency

of the $\text{CB}\ell\ell$ trigger was improved by the retuning, with efficiencies up to double those of the 2023 tuning in some modes. As anticipated, the efficiencies of the $\text{CB}\ell\ell$ trigger remain below those of the $\text{MVA}\ell\ell$ trigger. As the cuts on the topological variables now align more closely to those of the $\text{MVA}\ell\ell$ trigger preselection, it can also be expected that there is a greater overlap between the $\text{CB}\ell\ell$ and $\text{MVA}\ell\ell$ triggers. Surprisingly, the efficiencies in $B_s^0 \rightarrow \ell^+\ell^-$ are improved in both modes, despite the structure of the 2023 $\text{CB}\ell\ell$ trigger being configured for $b \rightarrow \ell\ell$ candidates. This is likely as the isMuon and electron acceptance cuts are significantly more efficient than the cuts required to suppress background in the 2023 $\text{CB}\ell\ell$ trigger, *i.e.*, those of the b candidate.




Figure 5.15: Combined trigger efficiencies of the 2023 (original) and 2024 (retuned) implementations of the $\text{CB}\ell\ell$ trigger, and the $\text{MVA}\ell\ell$ trigger as deployed in 2023, computed on signal MC for 7 different final states containing a e^+e^- or $\mu^+\mu^-$ pair.

Retuning of the topological b -hadron trigger for 2025 data-taking

As discussed in Section 3.2.3 and Section 5.1.2, LHCb has long employed a set of inclusive trigger lines to select b -hadron decay candidates based on their topological properties, referred to collectively in the topological b -hadron (b -topo.) trigger. Ahead of Run 3, the b -topo. trigger was redeveloped, moving from the boosted decision trees employed in Runs 1 and 2 [157] to a monotonic Lipschitz NN architecture [154, 158]. The Run 3 b -topo. trigger was developed outside of the work of this thesis; however, this is not the case for the retuning of this trigger for 2025 data-taking. The development, discussed in more detail in Ref. [158] is thus summarised here, whilst the retuning is discussed in more detail.

The b -topo. trigger in Run 3 consists of four HLT2 lines. These lines select two- and three-body candidates³, with and without a requirement that at least one of the final-state particles has a response in the muon system. Two monotonic Lipschitz NNs are employed, trained to classify two- and three-body candidates as either signal or background. These NNs take as input kinematic and topological features of the candidates and their subsequent particles, as listed in Table B.4 in Appendix B.2.

The monotonic Lipschitz NN architecture, laid out in Ref. [154], was chosen as it allows for monotonicity of the NN response to be enforced in specific features of choice, providing robustness against changing detector conditions. This monotonicity is achieved by constraining the gradient of the NN by the Lipschitz constant L , which is taken to be

³Three-body candidates are constructed as the combination of a two-body candidate with an additional track.

greater than 0, and for which the constraint is stronger for larger values. The effect of these constraints is shown for a toy system with two features, comparing the monotonic Lipschitz NN with an unconstrained NN and a monotonic BDT in Figure 5.16 [154].




Figure 5.16: MVA response across a toy space in features of the b -topo. trigger from Ref. [154].

The two- and three-body candidates are subject to a loose cut-based preselection before the NNs are applied. This preselection is intended to remove candidates which are kinematically or topology extremely unlikely to originate from real b -hadron decays. The cuts applied in this preselection are listed in Table B.5 in Appendix B.2. Requirements made on the final-state particles ensure suppression at low- p_T and low track quality, where the signal purity is very low, such that the number of possible two- and three-body combinations is maintained at a manageable level. Requirements made on the combinations significantly reduce the number of candidates to which the NNs must be applied, improving the throughput of the b -topo. trigger.

The NNs are trained using minimum bias MC as a background sample and MC from a mixture of specific decay channels as a signal sample. A custom truth-matching is applied to the minimum bias sample to remove any contributions originating from a b -quark. The signal sample is split equally between each of its component channels.

The NNs were trained by minimising a binary cross-entropy (BCE) loss:

$$\mathcal{L} = -\frac{1}{N} \sum_i [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)], \quad (5.6)$$

where y_i is the class label of each event (0 for background events and 1 for signal events) \hat{y}_i is the NN output [159].

The trained models are applied to the preselected candidates, with thresholds on the responses of the NNs defined such that the combined b -topo. trigger provides a given bandwidth of output. In the pre-2025 iteration of the b -topo. trigger, these thresholds were determined by processing HLT1 -filtered samples running instances of the trigger lines with different thresholds. The corresponding efficiencies and bandwidths of these are determined to establish a Pareto front, the set of thresholds which are most efficient for a given rate [160]. This method provides thresholds to the precision of the separation between each threshold value studied; however, this is limited by the computational resources required to run a large number of the trigger line instances at once. The allocation of bandwidth between the four lines, particularly between the standard and muonic pairs of lines, is somewhat subjective, and thus was determined by the expertise of the development team with input from analysts.

During 2024 data-taking, as a step in the validation of the data, the efficiencies of the b -topo. trigger, such as those shown in Fig. 5.17 were studied. These studies revealed two issues: the global efficiencies in each channel were lower than anticipated, *i.e.*, than in Fig. 3.14, and the efficiencies at high lifetime tail off rapidly above 5 ps. The former limits the statistics available to analyses, whilst the latter introduces problems in modelling decay time, as high lifetime events are essential to describing the cosh term in Eq. 2.32.

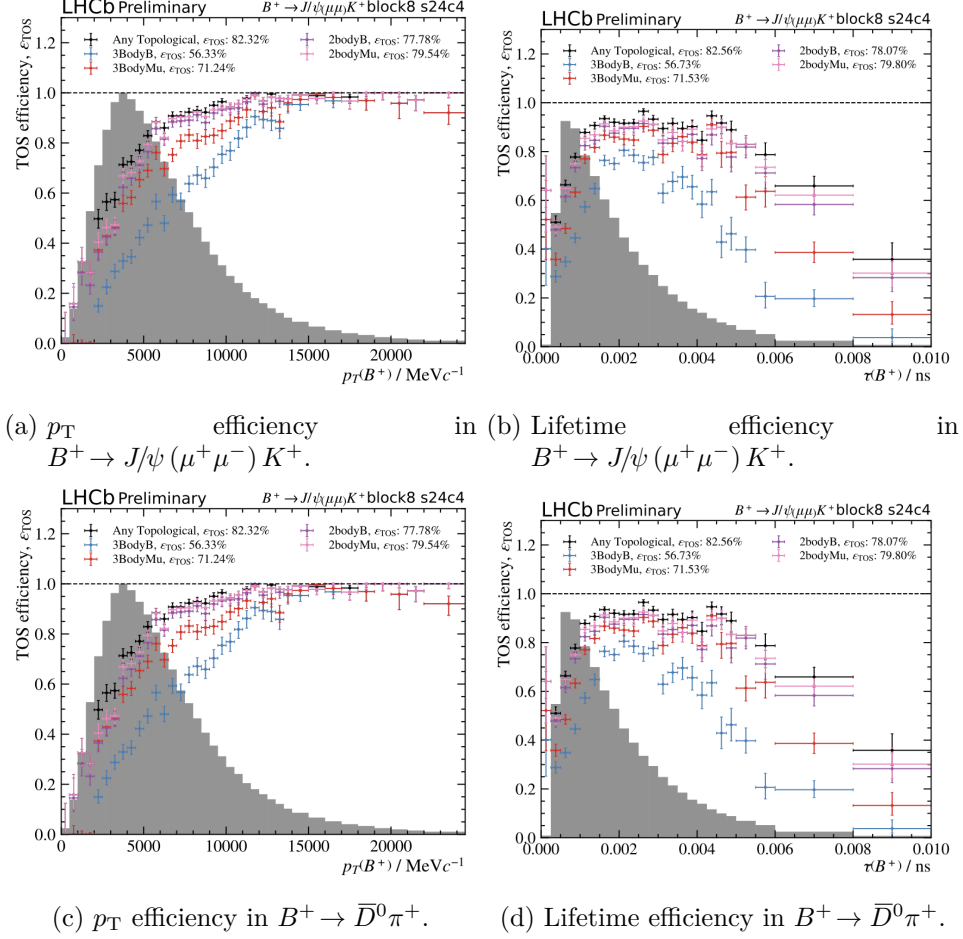


Figure 5.17: Efficiencies of the b -topo. trigger in late 2024 data, evaluated with the TISTOS method (described later in Section 5.2.2).

Improvement of the global efficiency is a difficult problem to solve, as the b -topo. trigger in Run 3 must contend with many more candidates than the Run 1/2 equivalent. This arises both from the removal of the L0 trigger enabling more soft candidates to enter HLT2 and from the higher pile-up in collisions in Run 3. To address this, the mixture of decay channels was revised. Additionally, requirements of monotonicity on the p_T of the 2-body combination and sum of p_T over its constituent final-state particles were removed, as were the equivalent requirements for the 3-body combination in the 3-body NN.

To understand the low efficiencies at high lifetime, the correlation of the input features of the NNs with the lifetimes of the candidates were studied in the signal and background samples. In most variables, no abnormal or unexpected correlations were seen; however,

when plotting events in the space of the minimum χ_{FD}^2 of the final state tracks vs two-body lifetime, as in Fig. 5.18, events are observed with a low minimum χ_{FD}^2 of final state tracks low for all lifetimes in the background sample. In contrast, no such distribution is seen in the signal sample. As a result of this, the signal purity decreases rapidly at high lifetimes, with the NNs consequently learning this behaviour and rejecting a disproportionate fraction of high lifetime events. The source of this background is believed to originate from combinations of tracks with a low χ_{FD}^2 originating from different PVs, such that these candidates appear to have high lifetimes.




Figure 5.18: Histograms in the space of minimum χ_{FD}^2 of final state tracks and two-body candidate lifetime in the background and signal training samples.

To mitigate the effect of this background on training, the distance correlation approach in Ref. [161], directly removes the correlation between the NN response and the lifetime by incorporating an explicit term in the loss. The BCE loss function in Eq. 5.6 is modified to include a penalty term to decorrelate the decay time, τ :

$$\mathcal{L} = -\frac{1}{N} \sum_i [y_i \log(\hat{y}_i) + (1 - y_i) \log(1 - \hat{y}_i)] + \lambda \cdot \text{dCor}(\hat{y}, \tau), \quad (5.7)$$

wherein λ is a tunable parameter and dCor is given by

$$\text{dCor}(\hat{y}, \tau) = \frac{\text{dCov}(\hat{y}, \tau)}{\sqrt{\text{dVar}(\hat{y}) \text{dVar}(\tau)}}, \quad (5.8)$$

in which \hat{y} is the NN response, $\text{dCov}^2(\hat{y}, \tau)$ is the distance covariance per Ref. [161] and $\text{dVar}(i)$ are the distance variances defined as $\text{dVar}(i) = \text{dCov}(i, i)$. The distance

covariance is defined such that it is 0 if its terms are independent. This modification propagates to the gradient of the loss as

$$\frac{\partial \mathcal{L}}{\partial \theta} = \frac{\partial \mathcal{L}_{\text{BCE}}}{\partial \theta} + \lambda \frac{\partial \text{dCor}(\hat{y}, \tau)}{\partial \theta}, \quad (5.9)$$

such that any correlation between the NN response and τ is penalised in the training. Whilst λ would typically require tuning to provide an optimal degree of penalisation, the exponentially decreasing nature of the background with decay time provides an inherent robustness to the scale of λ .

In the course of this retuning, the process by which the thresholds on the NN responses is determined was also consolidated. Whereas before the thresholds were determined by processing of HLT1 -filtered samples with many instances of the trigger across a range of thresholds, a refactoring of the *b*-topo. trigger separated the response of the NNs from the HLT2 lines. This enables the inclusion of NN responses in the resulting NTuples, allowing for samples to be processed with these thresholds set arbitrarily low and then artificially raised in a simple post-processing step. Not only does this reduce the time taken to optimise these threshold as only a single instance of the *b*-topo. trigger need be included in the processing, this also enables for a continuous optimisation of the thresholds, which in turn means that samples need not be reprocessed to determine threshold to a higher precision.

The consolidated optimisation procedure is as follows. The thresholds can be written as $\vec{c} = \{c_2, c_{2\mu}, c_3, c_{3\mu}\}$ for the 2-body, 2-body muonic, 3-body and 3-body muonic *b*-topo. trigger lines, respectively. As the thresholds are known to be below but very close to 1, these are parameterised as $\vec{c} = \vec{c}(\vec{x})$ wherein $\vec{x} = \{x_2, \delta x_{2\mu}, x_3, \delta x_{3\mu}\}$ to improve the stability of optimisation. For the thresholds of the standard lines, these are parameterised as

$$c_i = 1 - 10^{-x_i}, \quad (5.10)$$

whilst for the muonic lines, the thresholds, which are looser than those of the standard line by construction, are given by

$$c_{i\mu} = 1 - 10^{-(x_i - \delta x_i)}. \quad (5.11)$$

To quantify the combined rate and efficiencies of the *b*-topo. trigger for any set of parameters \vec{x} , MC samples of minimum bias and signal (specifically $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$, $B^0 \rightarrow D^{*-} \mu^+ \nu_\mu$, $B^0 \rightarrow D^- \pi^+$ and $B^+ \rightarrow \bar{D}^0 \pi^+$) were processed with no threshold applied to the NN responses, *i.e.*, with only the preselection applied. The response of the NNs on each event were then attached in the corresponding NTuples. If the rate of events passing the preselection, R_0 , is known (as can be computed from the HLT1 output rate), then the rate of the *b*-topo. trigger for a given \vec{x} can be computed as

$$\begin{aligned} r(\vec{c}(\vec{x})) &= R_0 \cdot f(\vec{c}(\vec{x}))_{\text{minbias}} \\ &= R_0 \cdot \frac{N(\vec{c}(\vec{x}))_{\text{minbias}}}{N_0}, \end{aligned} \quad (5.12)$$

wherein $f(\vec{c}(\vec{x}))_{\text{minbias}}$ is the fraction of minbias events passing the preselection of any *b*-topo. line (N_0) which also pass any of the *b*-topo. lines for NN thresholds $\vec{c}(\vec{x})$ ($N(\vec{c}(\vec{x}))$). As a proxy for the full efficiencies (*i.e.*, those accounting for the preselection cuts), the equivalent fraction in each channel i , $f(\vec{c}(\vec{x}))_i$, is used, dubbed the pseudo-efficiency. This can be used as the preselection cuts are independent of the threshold; the complete efficiency in channel i is highest when $f(\vec{c}(\vec{x}))_i$ is highest.

An optimal set of thresholds would give the highest possible efficiency in a given channel at a specified rate/bandwidth. The optimisation of these thresholds can therefore be constructed by maximising the efficiency over several decay channels, subject to a constraint of the output rate to a target rate R within some threshold α_R . Taking the total efficiency over several decay channels to be a weighted sum, in which the weights w_i encode the subjectivity of the bandwidth allocation between two- and three-body lines and standard and muonic lines, and the constraint on the rate to be Gaussian, the quantity to be maximised can be written as:

$$\ell_c = \exp \frac{(r(\vec{c}(\vec{x})) - R)^2}{2\alpha_R^2} \cdot \sum_i w_i \varepsilon_i(\vec{c}(\vec{x})). \quad (5.13)$$

Re-expressing this as a minimisation problem for numerical stability, *i.e.*, as $\mathcal{L}_c = -\log \ell_c$, the quantity to be minimised is

$$\mathcal{L}_c = \frac{(r(\vec{c}(\vec{x})) - R)^2}{2\alpha_R^2} - \log \sum_i w_i \varepsilon_i(\vec{c}(\vec{x})). \quad (5.14)$$

This minimisation was performed with `iminuit`, which provides a Python interface to the `Minuit2` minimiser [162, 163], for rate working points $R \in [6, 18]$ kHz in 200 Hz intervals, with a tolerance of $\alpha_R = 25$ Hz applied throughout. To improve the stability of the minimisation and to encourage monotonically decreasing (increasing) values of x_2 and x_3 ($x_{2\mu}$ and $x_{3\mu}$), the upper (lower) limit on each parameter for a given interval is taken from the respective value obtained for the previous interval. Additionally, the lower limits x_2 and x_3 are set to 90% of their values in the previous interval. This is necessary as the system is highly degenerate: local minima exist in the cases in which the majority of the bandwidth is allocated to the two-body or three-body lines alone. Discontinuities in the thresholds could arise if two consecutive minimisations reach different local minima. A similar degeneracy is also present between the standard and muonic lines. The system is biased towards specific local minima based on the chosen values of w_i : larger weights to channels with many final-state particles induce a greater allocation of bandwidth to the 3-body lines; larger weights to channels with muons naturally induce a greater allocation to the muonic lines. The weights chosen for this minimisation, listed in Table 5.1, were set to approximately balance these effects, with slight precedence given to the 2-body lines and to the standard lines as these provide coverage to a broader group of decays.

Table 5.1: Per channel weights applied in the minimisation of Eq. 5.14 for the optimisation of b -topo. trigger NN thresholds in 2025.

Channel	Weight, w_i
$B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$	1/3
$B^0 \rightarrow D^{*-} \mu^+ \nu_\mu$	2/3
$B^+ \rightarrow \bar{D}^0 \pi^+$	1
$B^0 \rightarrow D^- \pi^+$	2

The rate-efficiency curves of the optimal parameters in each of the optimisation channels⁴ are shown in Fig. 5.19. The pseudo-efficiencies rise monotonically with increases in the

⁴These differ from a Pareto front in that these are not necessarily the maximal efficiency with respect to the specific channel. Rather, these are the efficiencies in each channel for the parameters which return the maximal total weighted efficiency, *i.e.*, per Eq. 5.13.

output rate as expected, indicating that the optimisation does not move between the degenerate minima discussed above.

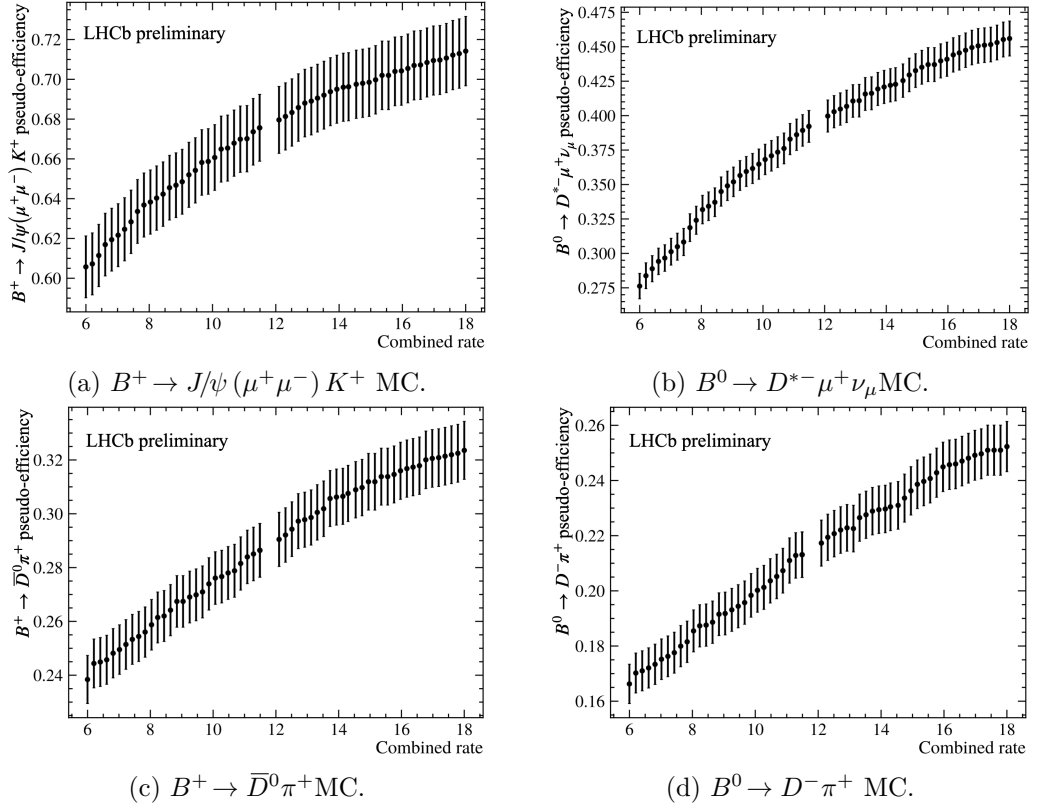


Figure 5.19: Efficiencies of the b -topo. trigger with optimised NN thresholds for rate working points $R \in [6, 18]$ kHz. Statistical uncertainties shown are fully correlated.

The NN thresholds set in 2025 data-taking, determined according to an expected output rate of 9 kHz, are determined to be 0.98858 for the 2-body line, 0.98205 for the 2-body muonic line, 0.99974 for the 3-body line and 0.99699 for the 3-body muonic line. Evaluating these in more detail through the automated bandwidth testing, these correspond to a total output bandwidth of the b -topo. trigger of ~ 8 GB/s.

Combined efficiencies in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$ and $B^0 \rightarrow D^-\pi^+$ were computed in MC simulated samples reflecting early 2024 data-taking for the 2024 and 2025 tunings of the b -topo. trigger and these updated thresholds. These efficiencies, calculated in the using `HltEfficiencyChecker`, in the space of the p_T and lifetime of spruced candidates, are shown in Fig. 5.20. A significant increase in efficiency for candidates with a high lifetime, particularly at low- p_T , is observed in both channels. Additionally, the global efficiency appears to be slightly improved in the 2025 tuning.

5.2.2 Estimating trigger efficiencies

Typically, efficiencies of selections in analyses are calculated by applying the selection to simulated signal, and taking the fraction of candidates passing the selection as the efficiency. This approach is valid, provided that the properties on which the selection is made are well-modelled in the simulated sample. This is not a generally true statement, in particular for complex selections such as those on trigger decisions, which typically depend on several selection cuts, including on the response of the particle identification

5.2 Development and performance of selections for the Run 3 software trigger

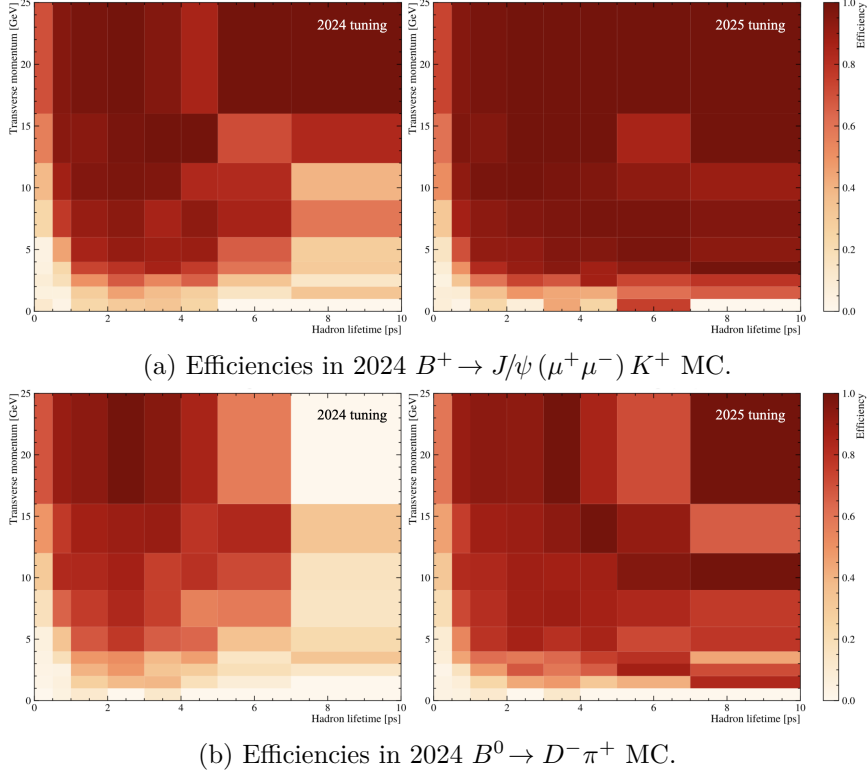


Figure 5.20: Efficiencies of the 2024 and 2025 b -topo. trigger configurations computed on 2024 MC samples.

systems, which are difficult to accurately represent in simulation. Therefore, another approach, discussed in detail in Ref. [164], is required to estimate the efficiency of trigger selections.

The TISTOS method was devised to provide a data-driven alternative to evaluating trigger efficiencies. Within this approach, events are categorised under a choice of a selection(s) of interest (*e.g.*, HLT1 or HLT2 lines) and signal of interest (*e.g.*, the B^+ in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$):

- Triggered on signal, TOS: the signal candidate within the event is sufficient to result in the line firing.
- Triggered independent of signal, TIS: another candidate in the event, independent of the signal candidate, is sufficient to result in the line firing.

These categories are not mutually exclusive, with events labelled both TIS and TOS henceforth referred to as TISTOS events. A further category, triggered on both, TOB, covers events which are neither TIS or TOS, wherein a combination of the signal candidate and another candidate is required for a line to fire. The TOB category is not discussed further as it forms a comparatively small proportion of events, estimated at approximately 0.5 % of events in b -decay.

Quantitatively, an event is labelled TOS if at least 70 % of detector hits in the event are also hits involved in the reconstruction of the signal candidate. An event is labelled as TIS if any reconstructed candidate in the event has fewer than 1 % of hits in common with the signal candidate. For composite particles, the hits from which these categories are evaluated are taken to be the pool of hits involved in any of the subsequent child

particles. For example, for a signal of the B^+ in the decay $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$, the pool of hits would be formed of any hits involved in the reconstruction of any of the μ^+ , μ^- and K^+ .

As described above, were all of the events, triggered or otherwise, present in the final sample, then the efficiency of a given trigger selection could simply be constructed, per Eq. 5.1, as

$$\varepsilon_{\text{Trig.}} = \frac{N_{\text{Trig.}}}{N_{\text{Tot.}}} \quad (5.15)$$

However, Eq. 5.15 can be rewritten to contain, explicitly, the events in the TIS category:

$$\varepsilon_{\text{Trig.}} = \frac{N_{\text{Trig.}}}{N_{\text{TIS}}} \frac{N_{\text{TIS}}}{N_{\text{Tot.}}} = \frac{N_{\text{Trig.}}}{N_{\text{TIS}}} \varepsilon_{\text{TIS}}, \quad (5.16)$$

wherein ε_{TIS} is the efficiency of requiring events to be categorised as TIS. This too must circumvent the requirement of $N_{\text{Tot.}}$, though this can be accomplished by measuring ε_{TIS} within the subsample of TOS events, *i.e.*,

$$\varepsilon_{\text{TIS|TOS}} = \frac{N_{\text{TISTOS}}}{N_{\text{TOS}}} \quad (5.17)$$

ε_{TOS} can be defined analogously, measuring the TOS efficiency within the subsample of TIS events. Eq. 5.17 can be substituted for ε_{TIS} in Eq. 5.16, since $\varepsilon_{\text{TIS}} \equiv \varepsilon_{\text{TIS|TOS}}$ within accountable kinematic correlations, such that $\varepsilon_{\text{Trig.}}$ is formed of measurable quantities:

$$\varepsilon_{\text{Trig.}} = \frac{N_{\text{Trig.}}}{N_{\text{TIS}}} \frac{N_{\text{TISTOS}}}{N_{\text{TOS}}}, \quad (5.18)$$

where $N_{\text{TIS}} N_{\text{TOS}} / N_{\text{TISTOS}}$ approximates $N_{\text{Tot.}}$.

This approach holds, provided that the correlations between the TIS and TOS subsamples are indeed small. The signal and independent candidates generally arise from each quark of the $q\bar{q}$ pairs produced in pp collisions, and thus are inherently correlated. Since trigger selections typically make extensive use of kinematic and topological information, these correlations can cause a bias in the trigger efficiency computed over all phase space. This bias is demonstrated in Fig. 5.21, where the transverse momentum of the B^+ candidate in simulated $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ events is compared with and without a TIS requirement at HLT1 .

To mitigate any bias from these correlations, the efficiency, or more specifically $N_{\text{Tot.}}$, can be calculated in bins of the phase space of the signal candidate, such that

$$\varepsilon_{\text{Trig.}} = \frac{N_{\text{Trig.}}}{N_{\text{Tot.}}} = \frac{N_{\text{Trig.}}}{\sum_i \frac{N_{\text{TIS}}^i N_{\text{TOS}}^i}{N_{\text{TISTOS}}^i}} \quad (5.19)$$

For sufficiently small phase space bins, the kinematic-dependent correlation and its subsequent bias becomes negligible. This is demonstrated in Fig. 5.22, wherein $\varepsilon_{\text{Trig.}}$ for the B^+ in simulated $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ decays is estimated using a variety of binning schemes, splitting the phase space in total and transverse momenta. As the binning becomes increasingly granular, the relative bias, *i.e.*, the difference between $\varepsilon_{\text{Trig.}}$ and the true efficiency, is decreased to below a 1%-level effect.




Figure 5.21: Comparison of $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ 2024 MC samples with no selection applied and with a TIS requirement from Ref. [164].

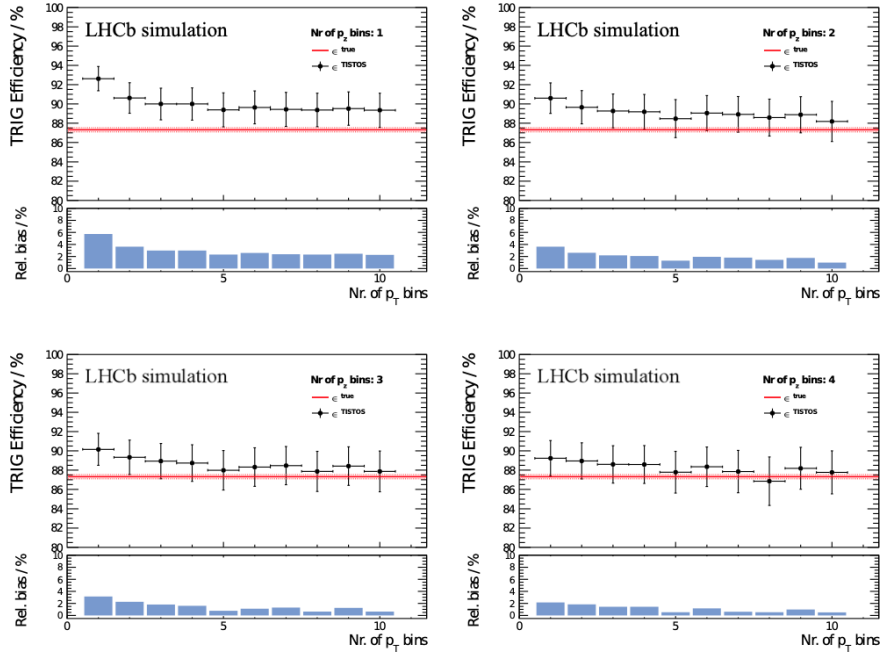


Figure 5.22: Comparison of $\varepsilon_{\text{Trig}}$, (per Eq. 5.19) with true efficiency in $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ Run 2 MC samples over increasingly granular phase space binnings from Ref. [165].

The TriggerCalib software package

The TriggerCalib software package, of which the author is the main developer, is a centralised implementation of the TISTOS efficiency calculations, providing a set of Python-based tools for use in LHCb physics analyses [164]. The package enables the evaluation of trigger efficiencies directly on data and is intended for use in high-statistics channels with well-understood backgrounds. These trigger efficiencies are stored as common Root [89] objects such that they can be seamlessly integrated into physics analyses and manipulated for purposes such as corrections to MC. Analysts can specify a binning scheme in which efficiencies are to be evaluated, with this being used both to provide local efficiencies across the binning scheme and the integrated efficiency, applying Eq. 5.19

to these local efficiencies.

TriggerCalib supports three methods of background mitigation, *i.e.*, subtraction or modelling of background contributions, each relying on a discriminating variable to separate signal from background, and thus to provide pure yields to Eq. 5.18.

The first, and simplest, method of background separation supported by TriggerCalib is sideband subtraction. The density of background, ρ , is estimated within one or more windows in the discriminating variable, *i.e.*,

$$\rho = \frac{\sum_i N_{\text{bkg.}}^i}{\sum_i \Delta_{\text{bkg.}}^i}, \quad (5.20)$$

for windows $\Delta_{\text{bkg.}}^i$ containing $N_{\text{bkg.}}^i$ events. The signal yield, or more specifically, the background-subtracted yield within some window $\Delta_{\text{sig.}}$, is then given by

$$N_{\text{sig.}} = N_{\text{Trig.}} - \rho \Delta_{\text{sig.}}. \quad (5.21)$$

Provided that the background is well-behaved, *i.e.*, approximately linear, within the signal window, the sideband windows can be chosen to appropriately estimate the background to be subtracted. However, this approach is too imprecise to be used for more complex background distributions, particularly those involving a peaking structure within the signal window.

The second method implemented in TriggerCalib is the fit-and-count method, as described in Section 5.1.2, in which an extended negative log-likelihood fit is performed to the discriminating variable in each of the bins and each of the yield categories (triggered, TIS, TOS and TISTOS). This exploits the understanding of background components to effectively and precisely separate between signal and background, and indeed between backgrounds if multiple backgrounds are present. Both the RooFit [166] and zfit [167] statistical analysis frameworks are supported for fitting, with the interface intentionally developed to be agnostic to the source of observable and PDF objects. However, performing likelihood fits in every bin can be time-consuming and computationally expensive, and ensuring that all fits converge can be difficult for finely-grained binning schemes.

The third and final method implemented in TriggerCalib is the *sPlot* method, in which a single likelihood fit per yield category is performed globally, and per-event *sWeights* are extracted from the resulting distributions [105]. This greatly reduces the number of fits which must be performed, from $4N$ to 4 for N bins, though this approach is only valid if the discriminating variable is independent of the control variables (or if any correlation cancels in the ratios between terms). Two tests of this, as prescribed by Ref. [168], are implemented in the package.

Firstly, a hypothesis test can be performed to confirm or reject whether the shape of the discriminating variable distribution is independent of the control variable(s). This test involves splitting the sample into two subsamples, defined as either side of a threshold in a control variable. Two fits to these subsamples are then constructed, using the model intended for use in computing *sWeights*: the first is a pair of independent fits, one to each subsample; the second is a simultaneous fit to both subsamples, with the parameters defining the shape of the distribution shared across the two subsamples. The corresponding likelihoods can be labelled as L_{H^1} and L_{H^0} , respectively, which correspond to the null hypothesis, H^0 , that the control and discriminating variables are independent,

and an alternate hypothesis where the variables are dependent, H^1 . A Q -statistic is defined from these likelihoods as

$$Q = -2 \cdot (\ln \sup\{L_{H^0}\} - \ln \sup\{L_{H^1}\}), \quad (5.22)$$

from which a p -value can be extracted from the χ^2 distribution with $N^{H^1} - N^{H^0}$ degrees of freedom, and used to reject or confirm H^0 .

Alternatively, and arguably more simplistically, the dependence of the discriminating and control variables can be tested by examining the Kendall rank correlation coefficient, τ [169]. Taking pure signal and background samples, *e.g.*, signal MC and a subsample of data from a sideband region, τ can be evaluated between the chosen control variable and discriminating variable in each sample. A p -value can be extracted against the null hypothesis that the variables are uncorrelated, *i.e.*, $\tau = 1$, and used to accept or reject this. The null hypothesis must be accepted in both the signal and background samples for the *sPlot* method to be valid.

It is not improbable that small correlations between the control and discriminating variables, cancel in the ratio of the efficiency, and thus do not result in a bias. This can be tested by comparing the resulting efficiencies from the *sPlot* method with those of the fit-and-count approach. It may be necessary to evaluate the latter in a more coarse binning scheme, to ensure that all bins are well-populated and that all fits converge.

The three background mitigation methods were evaluated on MC simulated samples of $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ decays with a toy combinatorial background injected, as shown in Fig. 5.23. The efficiencies returned by each method were consistent with one another, with exact consistency between the `Roofit` and `zfit` implementations. The uncertainties on efficiencies computed with the fit-and-count and *sPlot* methods are slightly larger than those computed with sideband subtraction, as these incorporate statistical uncertainties in the fitting, *i.e.*, from correlations between the floating parameters, as discussed in more detail in the next section.




Figure 5.23: Comparison of background mitigation methods from Ref. [164].

Along with carefully considering how trigger efficiencies should be evaluated, the statistical and systematic uncertainties must be similarly carefully considered.

Handling of statistical and systematic uncertainties

Turning our attention first to the statistical uncertainties, for a generic efficiency ε , with numerator n and denominator N , then a naïve propagation of Poisson uncertainties

would yield statistical uncertainties of

$$\sigma_\varepsilon = \varepsilon \sqrt{\frac{1}{n} + \frac{1}{N}}. \quad (5.23)$$

This breaks down in the two extrema of $n = 0$ ($\varepsilon = 0$) and $n = N$ ($\varepsilon = 1$) [170]. We must therefore turn to a more detailed treatment of the statistical uncertainties, or more specifically, of the interval at 68% confidence around the estimated efficiency. In particular, we look to the Wilson score interval introduced in Ref. [171].

The Wilson score interval considers true and measured efficiencies, ε and $\hat{\varepsilon}$, respectively. Under the approximation of the normal and binomial distributions, the half-width of the normal distribution at a confidence of $1 - \alpha$, z_α ($z_\alpha \approx 1$ for $\alpha = 0.68$) can be stated as

$$z_\alpha^2 \approx \frac{(\varepsilon - \hat{\varepsilon})^2}{\text{Var}(\hat{\varepsilon})}. \quad (5.24)$$

Incorporating the variance of the binomial distribution, which is given by

$$\text{Var}(\hat{\varepsilon}) = \frac{\hat{\varepsilon}(1 - \hat{\varepsilon})}{N}, \quad (5.25)$$

one can arrive at the intervals above/below $\hat{\varepsilon}$, ε_{\pm} , of

$$\varepsilon_{\pm} = \hat{\varepsilon} \pm \sigma_{\varepsilon, \pm} = \frac{1}{1 + \frac{z_\alpha^2}{N}} \left(\hat{\varepsilon} + \frac{z_\alpha^2}{2N} \pm \frac{z_\alpha}{N} \sqrt{N\hat{\varepsilon}(1 - \hat{\varepsilon}) + \frac{z_\alpha^2}{4}} \right). \quad (5.26)$$

This method provides coverage into the extrema from before, and, by considering the interval around $\hat{\varepsilon}$, results in asymmetric uncertainties which properly reflect the manner in which $\hat{\varepsilon}$ may fluctuate.

This approach still relies on the statistical uncertainties on the numerator and denominator being Poissonian; however, there are two cases relevant to the TISTOS method in which this is not strictly true:

1. The terms may contain non-Poissonian contributions, *e.g.*, through correlations in fitting, the uncertainties on the numerator/denominator may not be perfectly \sqrt{n}/\sqrt{N} .
2. In Eq. 5.19, the denominator is not a single counted term, so much as a combination of overlapping counted terms.

Considering the second point, we must take care to account for the overlap between the categories in the denominator of Eq. 5.19. Taking the approach laid out in Ref. [165], the denominator can be written as a combination of statistically independent terms:

$$\frac{N_{\text{TIS}}^i N_{\text{TOS}}^i}{N_{\text{TISTOS}}^i} = \frac{(\alpha^i + \gamma^i)(\beta^i + \gamma^i)}{(\gamma^i)^2}, \quad (5.27)$$

where $\alpha^i = N_{\text{TIS}}^i - N_{\text{TISTOS}}^i$, $\beta^i = N_{\text{TOS}}^i - N_{\text{TISTOS}}^i$ and $\gamma^i = N_{\text{TISTOS}}^i$. The statistical uncertainty on the denominator can then be written as

$$\sigma_{N_{\text{Tot.}}}^2 = \sum_i \sigma_{N_{\text{Tot.}}^i}^2 = \sum_i \left(\frac{\beta^i + \gamma^i}{\gamma^i} \right)^2 \sigma_{\alpha^i}^2 + \left(\frac{\alpha^i + \gamma^i}{\gamma^i} \right)^2 \sigma_{\beta^i}^2 + \left(1 - \frac{\alpha^i \beta^i}{\gamma^{i2}} \right)^2 \sigma_{\gamma^i}^2. \quad (5.28)$$

As to incorporating this and other non-Poissonian contributions to the statistical uncertainties on the yields, Ref. [172] provides a solution. The measurement of \hat{n} can be related to the true value n as

$$\hat{n} = n + z_1\sqrt{n} + z_2\sigma_b, \quad (5.29)$$

wherein z_i are separate random numbers with $E[z_i] = 0$, $E[z_i^2] = 1$ and σ_b^2 describes the contribution to the variance from non-Poissonian contributions. In this case, the variance on \hat{n} can be written as,

$$\text{Var}(\hat{n}) = n + \sigma_b^2, \quad (5.30)$$

where \hat{n} can be substituted for n , provided that \hat{n} is a sufficiently accurate estimator of n . Incorporating this in the variance of $\hat{\varepsilon}$ and recomputing the Wilson score interval, the generalised Wilson interval is, as per Ref. [172], given by

$$\varepsilon_{\pm} = \frac{1}{1 + \frac{z_{\alpha}^2}{N} \left(1 - \frac{\sigma_{n,b}^2 + \sigma_{N,b}^2}{N}\right)} \left(\hat{\varepsilon} + \frac{z_{\alpha}^2}{2N} \left(1 - 2\frac{\sigma_{n,b}^2}{N}\right) \pm \frac{z_{\alpha}}{N} \sqrt{\hat{\varepsilon}^2 \left(\sigma_{n,b}^2 + \sigma_{N,b}^2 - N\right) + \hat{\varepsilon} \left(N - 2\sigma_{n,b}^2\right) + \sigma_{n,b}^2 + \frac{z_{\alpha}^2}{4} \left(1 - 4\frac{\sigma_{n,b}^2 \sigma_{N,b}^2}{N^2}\right)} \right). \quad (5.31)$$

In the case of the efficiency described in Eq. 5.19, the corresponding non-Poissonian contribution to the variance of each yield can be described as

$$\sigma_{\text{Trig},b}^2 = \sigma_{\text{Trig}}^2 - N_{\text{Trig}}, \quad (5.32a)$$

$$\sigma_{\text{Tot},b}^2 = \sigma_{\text{Tot}}^2 - N_{\text{Tot}}, \quad (5.32b)$$

where all quantities except for σ_{Tot} are measured directly, and σ_{Tot} can be obtained via Eq. 5.28.

The systematic uncertainties on a trigger efficiency are inextricably linked to the channel of interest, and thus there is no concrete method by which to estimate them. However, the leading contributions to these systematic uncertainties are the choice of the tag (typically TIS) selection, the choice of phase-space variables and binning, and the choice of background mitigation.

To estimate the effect of each of these, analysts should estimate their efficiencies whilst varying the approach taken. In the case of the choice of TIS selection, an alternative selection can be used, and the resulting trigger efficiencies compared with the nominal trigger efficiencies. A similar approach can be taken for the choice of phase-space variables and binning. An alternate choice of variables can be applied, and the granularity of the bins varied, to obtain an alternative efficiency for comparison. Finally, for the choice of background mitigation method, the choice of alternate method is entirely analysis dependent, *i.e.*, the alternative method can only be applied if the channel of interest satisfies the requirements for the approach to be valid. For example, sideband subtraction would not be applicable for a channel with peaking background beneath the signal distribution in the discriminating variable; however, fit-and-count and the *sPlot* method could be compared, provided the discriminating and control variables are not strongly correlated.

5.2.3 Inclusive trigger efficiencies in 2024 and 2025 data-taking

With the selections of HLT1 and HLT2 developed, and in many cases redeveloped, the final step in commissioning these selections is to evaluate and understand their performance, namely their efficiencies, during data-taking. Of particular importance are the inclusive triggers at HLT1 and HLT2 which provide the majority of the selected events. The TOS efficiencies of the combined HLT1(Two)TrackMVA were evaluated in early 2024 data (reported in Ref. [173]), late 2024 data, and early 2025 data (reported in [174]). Similarly the TOS efficiencies of the combination of HLT1(Two)TrackMVA with the HLT2 b -topo. trigger were evaluated in late 2024 data and in early 2025 data in (reported in Ref. [174]). This section presents the approach by which these efficiencies were estimated and evaluates these against their equivalent efficiencies in Run 2.

All efficiencies from 2024 and 2025 data discussed in this section were computed using `TriggerCalib`, binning in the p_T of b -hadron candidates and applying sideband subtraction to remove background. The 1.2 fb^{-1} Block 1 dataset from Table B.1 in Appendix B of pp collision data taken at $\mu = 4.4$ is taken to represent early 2024 data-taking. The 0.74 fb^{-1} Block 7 dataset and 0.44 fb^{-1} Block 8 dataset from Table B.1 taken at $\mu = 5.3$ represent late 2024 data-taking for the HLT1 and the combined HLT1 and HLT2 efficiencies, respectively. The 1.74 fb^{-1} magnet-up 2025 dataset from Table B.1 taken at $\mu = 5.3$ represents early 2025 data-taking. Efficiencies with respect to muonic, electronic and hadronic final states are represented by the decays $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$, $B^+ \rightarrow J/\psi (e^+ e^-) K^+$ and $B^0 \rightarrow D^- \pi^+$, respectively, in which the efficiencies are studied.

Efficiencies of HLT1 were computed on NTuples produced from the direct output of exclusive HLT2 lines for each channel, using HLT1(Two)TrackMVA as both TIS and TOS. In each case, these are compared with the product of the efficiencies of the corresponding L0 trigger (`L0Muon/L0DiMuon` for muonic modes, `L0Electron` for electronic modes, `L0Hadron`) and HLT1(Two)TrackMVA in Run 2, as taken from Ref. [62].

Combined HLT1 and HLT2 efficiencies were computed from the output of Sprucing lines corresponding to each channel. For both TIS and TOS, HLT1(Two)TrackMVA was taken at HLT1; at HLT2 HLT2Topo{Two,Three}Body was taken. Similarly to the HLT1 efficiencies, these are compared with the products of the respective L0 trigger, HLT1 and HLT2 efficiencies in Run 2 taken from Ref. [62].

Trigger efficiencies in HLT1

The HLT1(Two)TrackMVA efficiencies are studied over the course of data-taking in 2024 and 2025. This provides an insight into the performance of the LHCb detector upgrades in concert, and comparisons of these efficiencies with one another demonstrate the relative improvements made to the trigger over this period. Firstly, the efficiencies of HLT1TwoTrackMVA in early 2024 data are shown, alongside their equivalent efficiencies incorporating the L0 trigger in Run 2, in Fig. 5.24. In muonic, electronic and hadronic b -hadron decay modes, the Run 3 trigger outperforms the Run 2 trigger, returning higher efficiencies in every p_T bin. In all three cases, the greatest gains are made for the softest b -hadron candidates, candidates which had previously had been discarded by the L0, which favours high E_T signatures.

In $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$, this improvement is smallest as the additional information from the muon stations enabled the `L0Muon` trigger to operate efficiently even at increased instantaneous luminosity, as shown in Fig. 5.4 in Section 5.1.2. The efficiency, particularly at low p_T , is nonetheless improved in muonic modes. In $B^+ \rightarrow J/\psi (e^+ e^-) K^+$, a

5.2 Development and performance of selections for the Run 3 software trigger

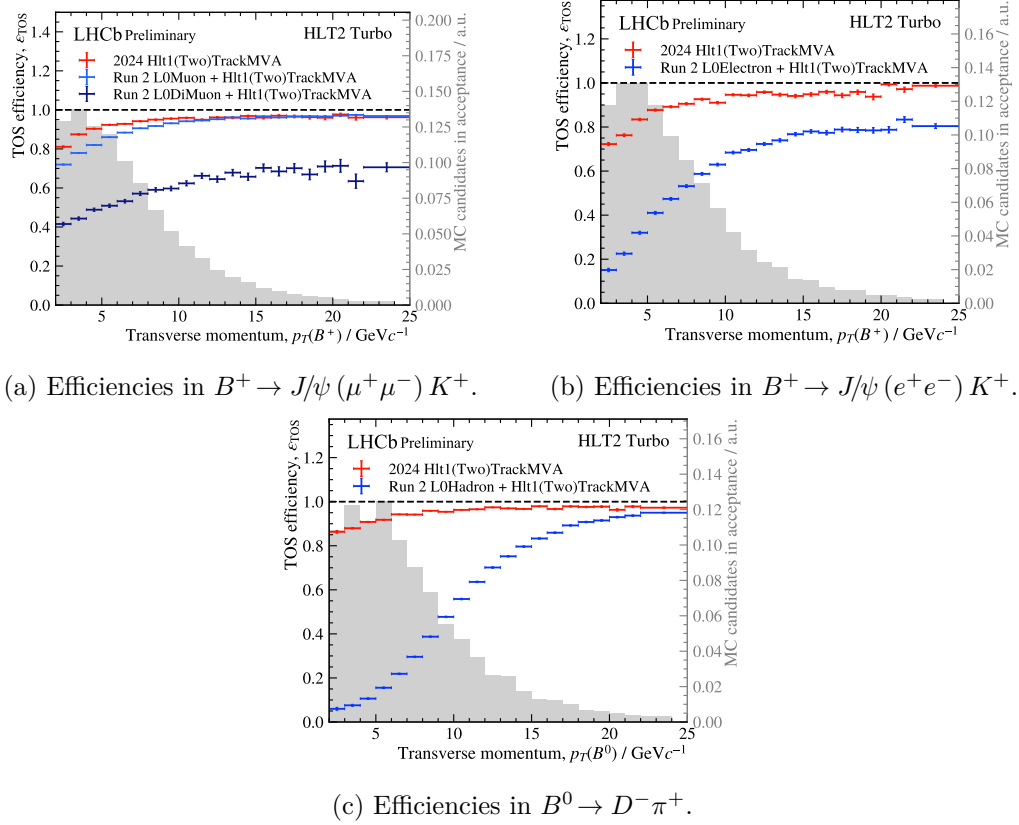
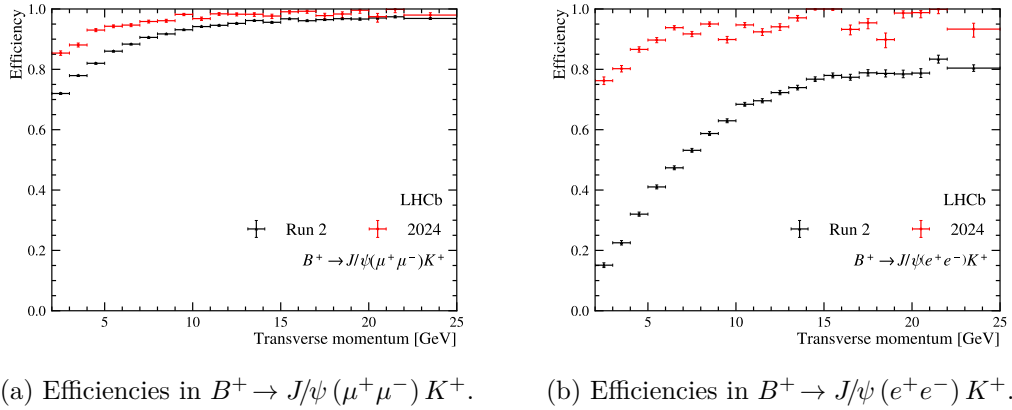


Figure 5.24: HLT1(Two)TrackMVA efficiencies for muonic, electronic and hadronic b -hadron decays in early 2024 data, shown in red. Equivalent efficiencies in Run 2 (the product of respective L0 trigger efficiencies with those of HLT1) are shown in blue. The distribution of b -hadron candidates produced in pp collisions which lie within the detector acceptance, taken from MC, is overlaid in grey. Figure is taken from Ref. [173].

much greater improvement in the efficiencies is observed. Again, the greatest improvements are seen at low p_T , with improvements of $\sim 4\times$ observed in the lowest p_T bins. The greatest increase in efficiency with respect to Run 3 is observed $B^0 \rightarrow D^-\pi^+$, where the efficiency in the lowest p_T bins is improved by $\sim 16\times$. Large improvements are seen even for b -hadron candidates with $p_T > 10 \text{ GeV}/c$. The improvement in the efficiency integrated over $p_T(B)$ is approximately $3\times$. Assuming that the relative improvement in HLT1 efficiency is the same for $B^0 \rightarrow D^-\pi^+$ is $B_s^0 \rightarrow D_s^-\pi^+$, this alone would improve the precision on $\langle A_{\text{untagged}}^s \rangle$ by $\sim 70\%$.

These efficiencies conclusively demonstrate that the Run 3 LHCb trigger can select events in real time at 30 MHz at significantly higher efficiencies than the Run 1/2 trigger. The upgrade of the trigger therefore fulfils the purpose of ensuring that LHCb can capitalise on the increase in instantaneous luminosity between Run 2 and Run 3.

Towards the end of 2024 data-taking, the event multiplicity was increased from $\mu = 4.4$ to $\mu = 5.3$. Busier events are more computationally demanding to process, which would typically require a change in trigger configuration to maintain the same throughput and output rate, *e.g.*, tighter selection requirements. However, the addition of a third GPU per-EB node in August 2024 provided an increase in the HLT1 throughput ahead of this increase in multiplicity and enabled a more performant reconstruction algorithm to be run. To understand the impact of these changes, the HLT1 efficiencies were evaluated for late 2024 data, as shown in Fig. 5.25.


 (a) Efficiencies in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$.


 (b) Efficiencies in $B^+ \rightarrow J/\psi(e^+e^-)K^+$.

 (c) Efficiencies in $B^0 \rightarrow D^-\pi^+$.

Figure 5.25: HLT1 (Two)TrackMVA efficiencies for b -hadron decays in late 2024 data (red). As in Fig. 5.24, the equivalent Run 2 L0 and HLT1 efficiencies (black) are overlaid.

Despite the increase in multiplicity, which normally would be expected to reduce the

trigger efficiency, the efficiencies in all three modes remain high, at comparable values to those shown in Fig. 5.24. In addition to demonstrating that HLT1 is capable of operating with a high efficiency at nominal data-taking conditions, *i.e.*, $\mu = 5.3$, this also shows that the trigger scales well with the addition of computing resources.

Between the end of 2024 data-taking and the recommencing of data-taking in 2025, improvements were made to the UT improving the hit efficiency. The additional UT hits can be used to provide a better fit of the tracks, particularly for softer tracks which are harder to determine due to their smaller curvature. The efficiencies of HLT1(Two)TrackMVA in each mode are once again evaluated, using samples of pp collisions from early 2025 ($\mu = 5.3$), shown in Fig. 5.26. As anticipated, the performance of HLT1 at low p_T is improved with respect to Fig. 5.25, further maximising the gains made by removing the L0 trigger.

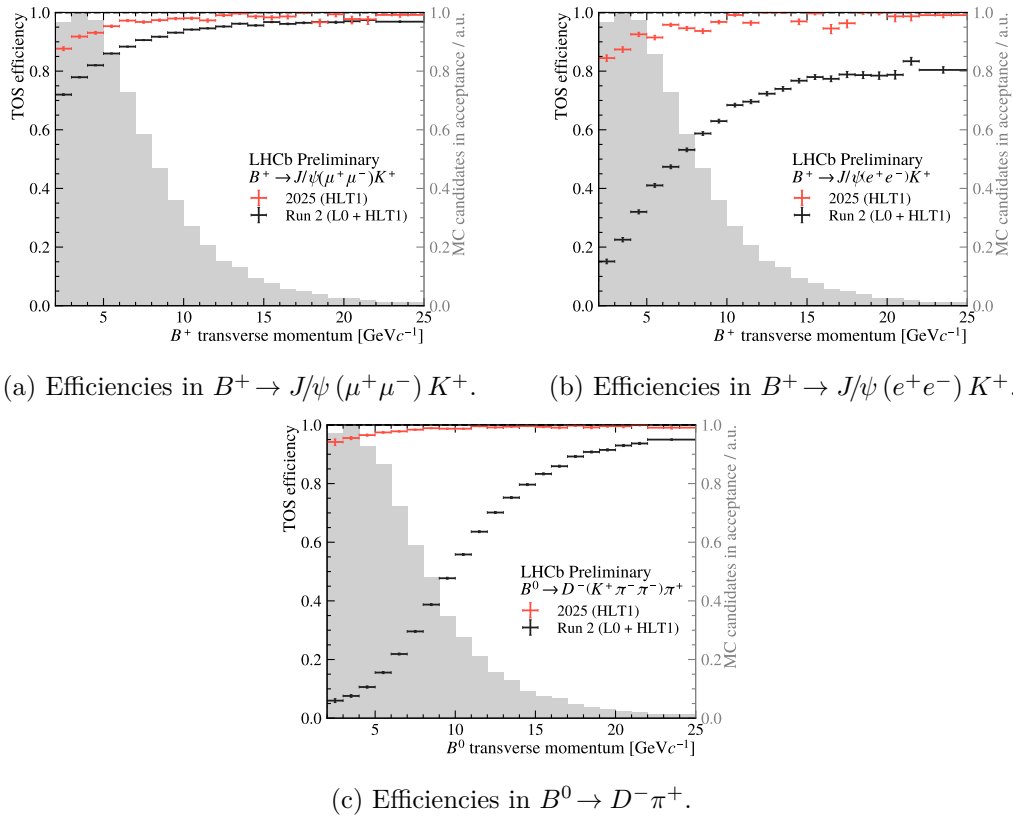



Figure 5.26: HLT1(Two)TrackMVA efficiencies for b -hadron decays in early 2025 data from Ref. [174]. As in Fig. 5.24, the equivalent Run 2 L0 and HLT1 efficiencies and p_T distribution of b -hadrons in the LHCb acceptance are overlaid.

Trigger efficiencies for HLT1 + HLT2

The combined HLT1(Two)TrackMVA and HLT2Topo{Two,Three}Body efficiencies are also studied in 2024 and 2025 data. These efficiencies can be used to understand the performance of the full trigger chain (HLT1, HLT2), to compare against the equivalent performance of the full Run 2 trigger chain (L0, HLT1, HLT2), and to quantify the contribution of the b -topo. trigger to the trigger efficiency.

The efficiencies of HLT1 and HLT2 together, evaluated in late 2024 data, are shown in Fig. 5.27. Whilst the efficiencies reported for 2024 are generally higher than their Run 2

equivalents, improvements at low b -hadron p_T are diluted by the efficiency of the b -topo. trigger. As described in Section 5.2.1, inefficiencies in the b -topo. trigger, both globally and locally at high candidate lifetimes, motivated the retuning of the b -topo. trigger early in 2025.


 (a) Efficiencies in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$.


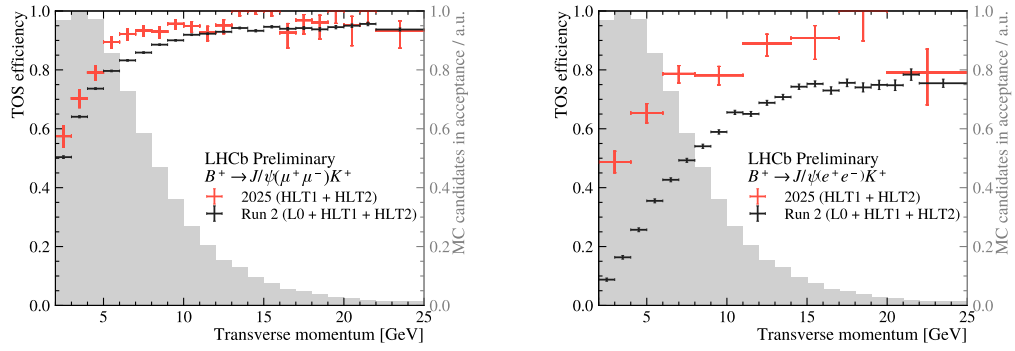
 (b) Efficiencies in $B^+ \rightarrow J/\psi(e^+e^-)K^+$.

 (c) Efficiencies in $B^0 \rightarrow D^-\pi^+$.

Figure 5.27: Combined HLT1(Two)TrackMVA and b -topo. trigger efficiencies for b -hadron decays in late 2024 data Equivalent efficiencies from Run 2 (the product of those of the L0, HLT1, and HLT2 triggers) are shown in black.

To confirm the improvements to the b -topo. trigger efficiencies brought about by the retuning described in Section 5.2.1, the equivalent efficiencies were evaluated in early 2025 data, as shown in Fig. 5.28. Increases in the trigger efficiency are observed in all three channels, with the greatest improvement seen in $B^0 \rightarrow D^-\pi^+$. The trigger efficiencies at high b -hadron p_T are generally very similar between late 2024 and early 2025; however, significant improvements are seen at low p_T . In particular, the efficiencies of the lowest p_T bins in $B^0 \rightarrow D^-\pi^+$ are improved by a factor of almost 2. Whilst these improvements are notable, a dilution of the improvements made by removing the L0 trigger is still present; work on identifying and mitigating the sources of inefficiencies at low p_T , in the b -topo. is ongoing at the time of writing.


From the efficiencies reported above, it is concluded that the retuning of the b -topo. trigger successfully improved the b -topo. trigger efficiencies across the LHCb physics programme in general and in hadronic modes in particular. These increases are in addition to the improvements to the efficiencies at high-lifetime discussed in Section 5.2.1.

5.2 Development and performance of selections for the Run 3 software trigger



(a) Efficiencies in $B^+ \rightarrow J/\psi(\mu^+\mu^-)K^+$.

(b) Efficiencies in $B^+ \rightarrow J/\psi(e^+e^-)K^+$.



(c) Efficiencies in $B^0 \rightarrow D^-\pi^+$.

Figure 5.28: Combined HLT1(Two)TrackMVA and b -topo. trigger efficiencies for b -hadron decays in 2025 data from Ref. [174]. Corresponding Run 2 efficiencies and b -hadron p_T distributions are overlaid per Figs. 5.27 & 5.24, respectively.

5.3 Conclusion

The upgrade of the LHCb experiment saw sweeping changes across the detector and trigger, centred upon the removal of the L0 trigger. The real-time approach applied to the trigger was successfully implemented and has demonstrated that it is capable of processing data at the 30 MHz pp collision rate. Beyond simply being capable, the results produced in the course of this work show clearly that the Run 3 trigger has realised, and in many cases exceeded, the anticipated trigger efficiency improvements [175].

Measurements performed with the Run 3 LHCb data, particularly statistically-limited measurements, stand to gain from the improvements in trigger efficiencies, *e.g.*, an improvement of $\sim 70\%$ in the $\langle A_{\text{untagged}}^s \rangle$ observable from these improvements alone as described above. However, these improvements can only be realised if these data, and the effects of the upgraded detector and trigger within, are studied and understood.

6 Measurement of $r_{J/\psi}$ in 2024 LHCb data

With the upgraded detector and trigger commissioned, it is crucial to demonstrate that LHCb can perform high-quality physics analyses with the recorded data. To this end, measurements are performed with this data, aiming to understand detection, reconstruction and selection effects in the data and how these are reflected in equivalent MC simulated samples, and to validate that the approaches and techniques intended for Run 3 analyses are indeed applicable.

The author's contributions have been primarily to the data-MC corrections, particularly those of the trigger efficiencies. As this measurement is ongoing, this chapter reflects the status of the measurement at time of writing.

6.1 Analysis strategy, samples and selections

This chapter discusses a measurement of the ratio between the $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ and $B^+ \rightarrow J/\psi (e^+ e^-) K^+$ integrated decay rates, known as $r_{J/\psi}$, in 2024 LHCb data, of which the author is a proponent. This quantity is a precursor to the equivalent quantity for the rare $B^+ \rightarrow K^+ \ell^+ \ell^-$ modes, R_K (discussed in Section 2.4.2) both directly and indirectly.

Directly, R_K is typically constructed as a double ratio for a given ratio in the dilepton invariant mass, $q^2 \in [q_a^2, q_b^2]$, *e.g.*, as in Ref. [77, 78]:

$$\begin{aligned}
 R_K(q_a^2, q_b^2) &= \frac{\int_{q_a^2}^{q_b^2} \frac{d\Gamma(B^+ \rightarrow K^+ \mu^+ \mu^-)}{dq} dq^2}{\int_{q_a^2}^{q_b^2} \frac{d\Gamma(B^+ \rightarrow K^+ e^+ e^-)}{dq} dq^2} \\
 &= \frac{\frac{N}{\varepsilon}(B^+ \rightarrow K^+ \mu^+ \mu^-)}{\frac{N}{\varepsilon}(B^+ \rightarrow K^+ e^+ e^-)} \bigg/ \underbrace{\frac{\frac{N}{\varepsilon}(B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+)}{\frac{N}{\varepsilon}(B^+ \rightarrow J/\psi (e^+ e^-) K^+)}}_{r_{J/\psi}}, \quad (6.1)
 \end{aligned}$$

wherein N/ε is the yield of a given decay, corrected by dividing by the corresponding efficiency to detect, reconstruct and select the decay. This is used as many of the systematic uncertainties involved in measuring the efficiencies cancel in the double ratio, reducing the scale of systematics from $\sim 5\%$ to below 1%.

Indirectly, $r_{J/\psi}$ provides a high-statistics proxy in which to understand effects such as offline selections and efficiency corrections. This relies on the fact the decays of $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ proceed so readily in the SM that any small NP contribution would have a negligible effect on $r_{J/\psi}$. It can thus safely be assumed to be 1; any observed deviation from this would indicate issues with the analysis techniques, rather than the presence of NP. It is on this same basis that the measurement described here can be used to validate the analysis techniques which would be applied to the rare modes in an equivalent measurement of $r_{J/\psi}$.

As the dimuon mode, which lies in the numerator of $r_{J/\psi}$, is more readily reconstructed and selected, it can also be experimentally advantageous to instead measure the reciprocal, $r_{J/\psi}^{-1}$:

$$r_{J/\psi}^{-1} = \frac{N(B^+ \rightarrow J/\psi (e^+e^-) K^+)}{N(B^+ \rightarrow J/\psi (\mu^+\mu^-) K^+)} \bigg/ \frac{\varepsilon(B^+ \rightarrow J/\psi (e^+e^-) K^+)}{\varepsilon(B^+ \rightarrow J/\psi (\mu^+\mu^-) K^+)}. \quad (6.2)$$

The former are extracted from pp collision data collected in 2024 and subject to a set of offline selections (described in Section 6.1.3) by performing negative log-likelihood fits to the B^+ mass, as laid out in Section 6.2. The latter are determined from MC simulation, correcting for differences between data the simulated samples. Determining these data-MC corrections is non-trivial and is thus described in detail in Section 6.3.

6.1.1 Data samples and trigger selections

This measurement examines the 2024 LHCb pp collision data from blocks 1 and 5-8, as listed in Table B.1 in Appendix B, which correspond to a total integrated luminosity of $\sim 4.5 \text{ fb}^{-1}$. Blocks 2-4 were not considered for this measurement as the trigger and reconstruction configurations were significantly different to the more stable configurations of the other blocks.

At HLT1, either HLT1TrackMVA or HLT1TwoTrackMVA, is required to trigger on (TOS) the B^+ candidate. At HLT2, candidates in the dielectron and dimuon modes are required to pass H1t2RD_BuToKpEE and H1t2RD_BuToHpMuMu_Incl_Bp, respectively. The cut-based selection requirements applied in each line are listed in Table C.1 & C.2 in Appendix C.1. The data samples are drawn from the output of these HLT2 lines, which record events to the Turbo stream.

Candidates are reconstructed by forming a good quality vertex (the J/ψ) from a pair of long tracks subject to loose PID requirements identifying the tracks as either muons or electrons. The properties of reconstructed electrons are highly dependent on the amount of bremsstrahlung emitted between their production and detection. Electrons are therefore categorised based on whether the electron has associated recovered bremsstrahlung photons (described in Section 3.2.2). An additional long track is attached to form another vertex (the B^+), where in H1t2RD_BuToKpEE this track is subject to loose PID requirements identifying the track as a kaon. In H1t2RD_BuToHpMuMu_Incl_Bp, no such loose PID requirement is made as the trigger line was developed to select $B^+ \rightarrow h^+ \ell^+ \ell^-$ decays; a requirement is instead applied in the offline selection described in Section 6.1.3

To improve the resolution of the B^+ invariant mass, DTF is applied to compute the B^+ mass from the final state particle kinematics within the constraints of the decay topology [111]. This is particularly important in the dielectron mode, where the B^+ invariant mass is significantly broader due to the emission of bremsstrahlung.

MC simulated samples of $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ and $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^{*0}$ decays used in the evaluation of efficiencies and determination of templates in the invariant mass fit were produced for this analysis, and are subject to the same online data processing as the data samples.

6.1.2 Background contributions

The background contributions in reconstructing and selecting $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ decays arise through similar mechanisms to those described in Section 4.1.2. Combinatorial background arises both from pairs of unrelated tracks to form the J/ψ candidate, and from additional unrelated tracks attached to form the B^+ candidate. The primary source of background from misidentified final-state particles is of $B^+ \rightarrow J/\psi (\ell^+ \ell^-) \pi^+$ decays, which are CKM-suppressed with respect to $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ and can be suppressed to a negligible level by imposing PID requirements on the associated hadron. Decays of $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^{*0}$, which proceed at a similar rate to $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$, induce a significant partially reconstructed background through candidates reconstructed missing the π^- from the K^{*0} .

Unlike in Section 4.1.2, the presence of leptons in the final state introduce two forms of background not discussed thus far. The first is the presence of semileptonic b -decays, dominated by $B \rightarrow K \ell \nu X$ decays, which contribute a K^\pm and ℓ^\mp to the candidate. These candidates present in the dilepton invariant mass as broad structures centred at ~ 2000 MeV: in the narrow dimuon invariant mass, a mass window can remove this; in the much broader dielectron mass, it is only possible to remove this background with a veto. The second of these backgrounds arises from the misidentification of the π^+ of a $D^0 \rightarrow K^- \pi^+$ decay as a ℓ^+ , with the $K^\pm \ell^\mp$ then contributing similarly to the semileptonic background. Finally, sub-tracks or copies of tracks, known as clones may cause a track to be considered twice in the reconstruction [176], *i.e.*, as the ℓ^\pm and K^+ , though this can be mitigated by requiring the opening angle between these particles to be above a minimum value.

The presence of duplicate candidates in each event is assumed to be suppressed by the selection described in the following chapter and thus no direct treatment is applied.

6.1.3 Offline selection

The background contributions discussed in the previous section are suppressed by applying cut-based selection requirements to the kinematic, topological and PID properties of the final state and composite particles. These selection requirements are summarised in Tab 6.1 and discussed in detail below. Selections are kept as similar as possible between the two modes, such that many of the effects from selection efficiencies cancel in Eq. 6.1. These selections also seek to reduce differences in the selections applied by the two HLT2 lines. As far as possible, the selections are also defined independent of the J/ψ (except for requirements on q^2), such that the selections would also be applied for the rare $B^+ \rightarrow K^+ \ell^+ \ell^-$ modes.

Firstly, candidates are required to have a dilepton invariant mass (q^2) is required to be compatible with the J/ψ invariant mass. In the dimuon mode, the J/ψ invariant mass peak is narrow and centred on the measured J/ψ mass; a requirement of $\sqrt{q^2 - 3.0969 \text{ GeV}^2} < 0.150 \text{ GeV}$ is applied to these candidates. In the dielectron mode, the J/ψ invariant mass peak is much broader and features a long radiative left tail; a window of $q^2 \in [6, 11] \text{ GeV}^2$ is applied accordingly.

Requirements on metrics of particle identity are applied for both leptons and the K^+ , all of which are required to have corresponding hits in RICH1/RICH2. The former are subject to requirements that electron candidates are within the ECAL acceptance and that muon candidates have corresponding hits in the muon stations, in addition to requirements on PID_e and PID_μ to suppress particle misidentification (primarily from

Table 6.1: Summary of selection requirements.

Decay mode	Selection	Cut applied
Both	Trigger (HLT1)	HLT1(Two)TrackMVA TOS on B^+
	Particle identity	$\text{PID}_k > 4$ and $\text{ProbNNk} > 0.2$ for K^+
	Track quality	$\text{hasRich} = 1$ for K^+ , ℓ_1 and ℓ_2
	Final state particle kinematics	$\chi_{\text{IP}}^2 > 9$ for K^+ , ℓ_1 , ℓ_2 and J/ψ
		$p_{\text{T}}(\ell) > 500 \text{ MeV}/c$ for ℓ_1 and ℓ_2
		$\min\{p_{\text{T}}(\ell_1), p_{\text{T}}(\ell_2)\} > 3 \text{ GeV}/c$
	Semileptonic veto	$p_{\text{T}}(K^+) > 300 \text{ MeV}/c$
Clone veto	$p(K^+) < 2 \text{ GeV}/c$	
$J/\psi \rightarrow e^+e^-$	Trigger (HLT2)	$m(K^{\pm}\ell^{\mp}) > 1880 \text{ MeV}/c^2$
	q^2 window	$K-\ell$ opening angle $> 1 \text{ mrad}$ for ℓ_1 and ℓ_2
	Electron identity	H1t2RD_BuToKpEE
	Charm veto	$q^2 \in [6, 11] \text{ GeV}^2$
$J/\psi \rightarrow \mu^+\mu^-$	Trigger (HLT2)	$\text{PID}_e > 2$ for e_1 and e_2
	q^2 window	inEcal
	Muon identity	$ m(K^{\pm}e^{\mp} [\text{as } \pi^{\mp}]) - m_{D^0} < 30 \text{ MeV}/c^2$ or $\text{HasBremAdded}(e^{\mp}) = 0$ and $\text{PID}_e > 6$ for
	Charm veto	$\text{HasBremAdded}(e^{\mp}) = 0$ and $\text{PID}_e > 6$ for
$J/\psi \rightarrow \mu^+\mu^-$	Trigger (HLT2)	H1t2RD_BuToHpMuMu_Incl_Bp
	q^2 window	$\sqrt{q^2} - 3.0969 \text{ GeV} < 0.150 \text{ GeV}$
	Muon identity	$\text{PID}_{\mu} > -2$ for μ_1 and μ_2
	Charm veto	isMuon
$J/\psi \rightarrow \mu^+\mu^-$	Charm veto	$ m(K^{\pm}e^{\mp} [\text{as } \pi^{\mp}]) - m_{D^0} < 30 \text{ MeV}/c^2$ or $\text{PID}_{\mu} > 6$ for μ^-
	Charm veto	$\text{PID}_{\mu} > 6$ for μ^-

pions). The latter are subject to selection cuts on PID_k and ProbNNk , which suppress backgrounds arising from misidentification of the kaon such as $B^+ \rightarrow J/\psi (\ell^+\ell^-) \pi^+$.

The final state particles are also subject to requirements on kinematic and topological variables, with the primary aim of aligning the selections on these quantities already applied at HLT2.

In addition to the global selection requirements discussed above, three vetoes are applied in this measurement. To remove contributions from $B \rightarrow K\ell\nu X$ decays, the invariant mass of the K^+e^- combination is required to be above $1880 \text{ MeV}/c^2$; no equivalent requirement is made in the dimuon mode as the $K^+\mu^-$ invariant mass distribution does not extend down to $1880 \text{ MeV}/c^2$. Contributions from misidentified D^0 decays are suppressed by applying stricter particle identification requirements to candidates for which the $K^+\ell^-$ [as π^-], where ℓ^- [as π^-] is the respective lepton with the pion mass hypothesis assigned, invariant mass falls within $30 \text{ MeV}/c^2$ of the measured D^0 mass. In the dielectron mode, this requirement is that the electron of the K^+e^- pair must have no associated recovered bremsstrahlung and that it has $\text{PID}_e > 6$; in the dimuon mode, this requirement is that the muon of the $K^+\mu^-$ pair has $\text{PID}_{\mu} > 6$. Finally, clone tracks are removed by enforcing that the opening angle between each lepton and the K^+ is at least 1 mrad radians.

6.2 Extraction of yields

The yields in each mode are extracted by means of negative log-likelihood fits to the B^+ mass as computed with DTF. Models of the B^+ mass in each mode are constructed as the sum of three components: signal, and the combinatorial and partially reconstructed $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^{*0}$ backgrounds discussed in Section 6.1.2.

The signal contribution is modelled by a single DSCB (Eq. A.1) in the dielectron mode and the sum of a DSCB and Gaussian distribution in the dimuon mode, labelled as $f_{\text{sig.}}$. These are fit to corresponding signal MC samples, as shown for Block 8 in Fig. 6.1. In the dielectron mode, the DSCB models the MC reasonably well, particularly in the long radiative left tail. In the dimuon mode, the greater statistics require the additional degrees of freedom afforded by the Gaussian component to provide a similarly suitable description of the MC.

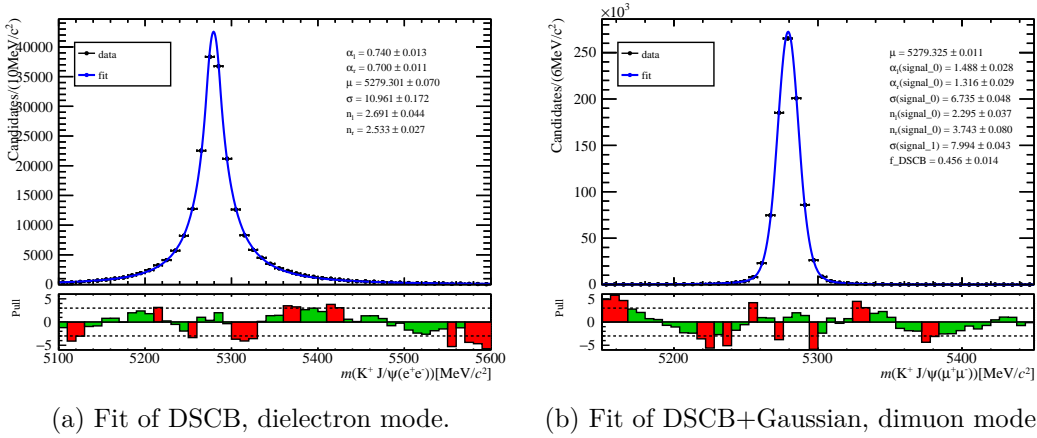


Figure 6.1: Fits of signal template PDFs to respective $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ MC in each decay mode.

The partially reconstructed background contributions cannot be modelled well analytically, as their resulting mass distributions are non-trivially related to the kinematics of the missing pion. This component is therefore modelled by a one-dimensional kernel density estimator [177], implemented using Gaussian kernels as `RooKeysPdf` in `RooFit` [166]. This empirical model, labelled f_{PR} , is fit to the $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^{*0}$ MC samples in each mode, as shown for Block 8 in Fig. 6.2.

The combinatorial background is modelled by a single exponential distribution, labelled $f_{\text{comb.}}$. The initial values of ℓ in the fits to data are determined by fitting the combinatorial model to candidates in the upper $m(B^+)$ sideband: $m(B^+) \in [5600, 6000]$ MeV/c^2 and $m(B^+) \in [5350, 6000]$ MeV/c^2 in the dielectron and dimuon datasets, respectively. These ranges are sufficiently far from the signal distribution in each mode that only the combinatorial contribution should be present. The combinatorial fits for Block 8 are shown in Fig. 6.3.

From these three components, the full invariant mass model in each mode is defined as

$$f_{\text{tot.}}(B^+) = N_{\text{sig.}} f_{\text{sig.}}(B^+) + N_{\text{comb.}} f_{\text{comb.}}(B^+) + N_{\text{PR}} f_{\text{PR}}(B^+), \quad (6.3)$$

where N_i describes the yield of each component. The shape parameters of each distribution are fixed, except for the combinatorial exponent, which is allowed to float and starts

6 Measurement of $r_{J/\psi}$ in 2024 LHCb data

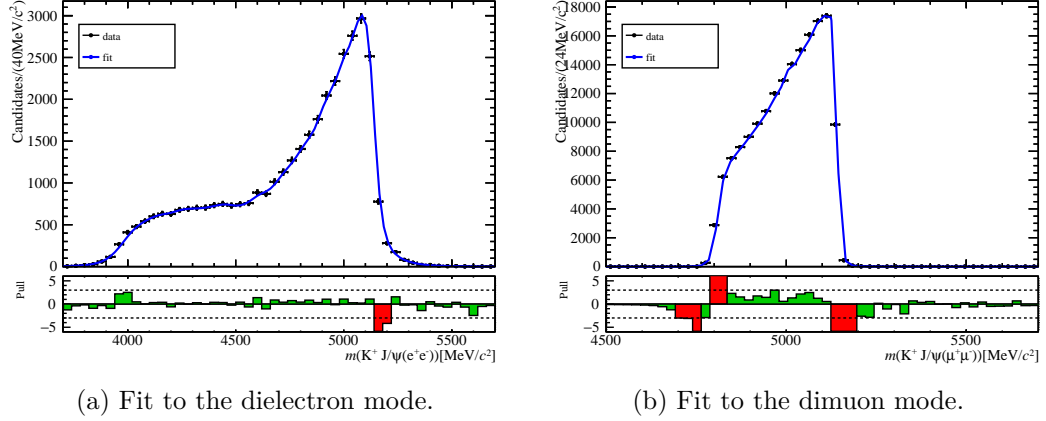


Figure 6.2: Fits of RooKeysPdf to partially reconstructed $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^{*0}$ MC samples in each decay mode.

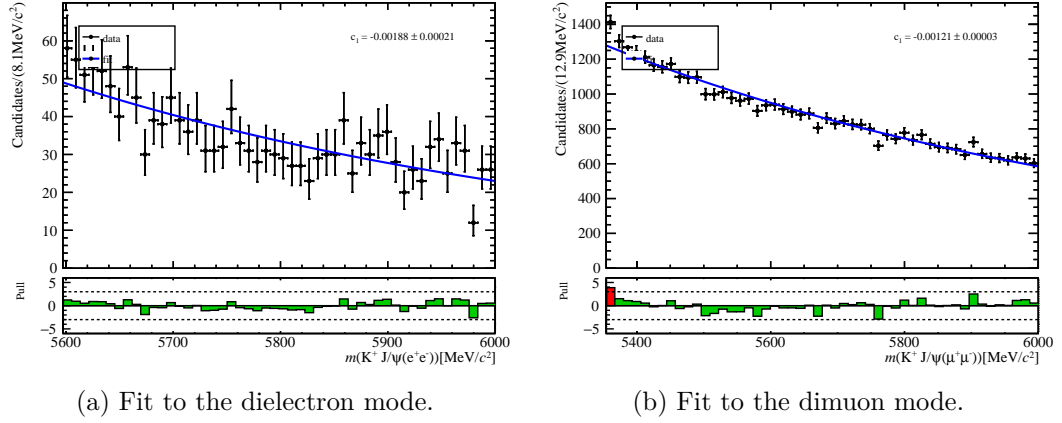


Figure 6.3: Fits of exponential distributions to the upper $m(B^+)$ sideband in data samples in each decay mode.

from its value determined in the fit to candidates above the signal peak in $m(B^+)$. To account for differences between the distribution of the B^+ mass in data and MC, namely the position and resolution of peaking structures, the mean and width parameters of f_{sig} are modified in the fit to data by floating additive and multiplicative terms, respectively, *i.e.*, μ_{MC} transformed as in Eq. 4.13a in Section 4.3.1 and the width parameters modified as

$$\sigma_{\text{data}} = \sigma_{\text{MC}} \cdot f_{\sigma}. \quad (6.4)$$

The mass models are fit to the B^+ mass over $m(B^+) \in [5100, 6000]$ MeV/ c^2 in each dilepton mode and data block independently, as shown in Fig. 6.4 for Block 8 and Fig. C.1 in Appendix C.2 for the other blocks. The yields extracted from each fit, and their ratios, are listed in Table 6.2.

All fits converge and describe the data well. The combinatorial and partially reconstructed backgrounds are difficult to disentangle as they overlap significantly at low $m(B^+)$; however, this interplay appears not to affect the description of the signal significantly. The ratios of the dielectron and dimuon yields in each component are fairly stable, with the ratio of signal yields varying by only $\sim 10\%$ between the smallest and greatest values.

6.3 Estimation of efficiencies and data-MC corrections

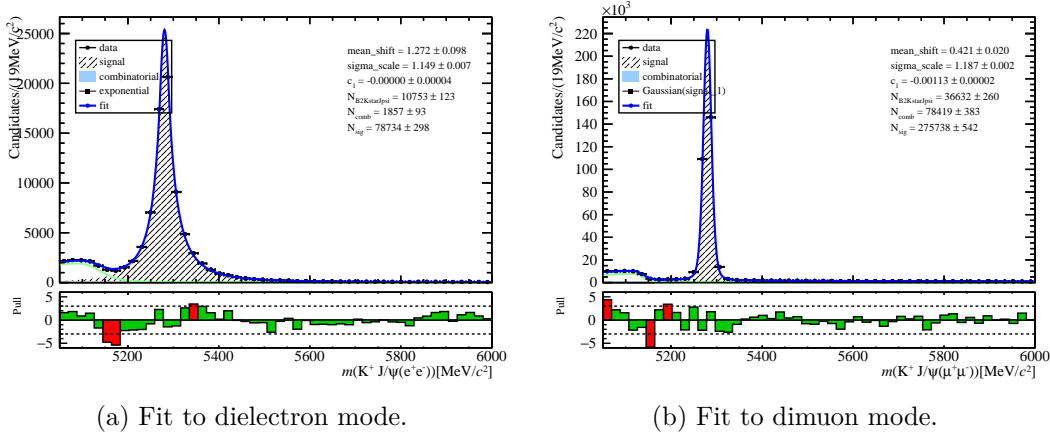


Figure 6.4: Fits to the B^+ DTF mass in Block 8 data samples. Corresponding yields are listed in Table 6.2. Equivalent fits in data blocks 1 and 5-7 are shown in Fig. C.1.

6.3 Estimation of efficiencies and data-MC corrections

The efficiencies in each mode can initially be computed as the product of the geometric acceptance efficiency (the proportion of generated particles which fall within the LHCb detector acceptance), $\varepsilon_{\text{acc.}}(\ell\ell) = N_{\text{acc.}}(\ell\ell) / N_{\text{gen.}}(\ell\ell)$, and the efficiency to reconstruct and select a $B^+ \rightarrow J/\psi(\ell^+\ell^-)K^+$ candidate (taken as the ratio of reconstructed and selected candidates to generated events in the LHCb acceptance), $\varepsilon_{\text{sel.}}(\ell\ell) = N_{\text{sel.}}(\ell\ell) / N_{\text{acc.}}(\ell\ell)$. As these are entirely dependent on MC simulated samples, the effect of differences between data and MC upon the efficiencies must be understood and accounted for.

This measurement must correct for data-MC differences in the q^2 resolution (in the dielectron mode), the efficiencies of track reconstruction, and the PID and trigger responses, the final state particle kinematics and the detector occupancy. The order in which these corrections are computed and applied is shown schematically in Fig. 6.5.

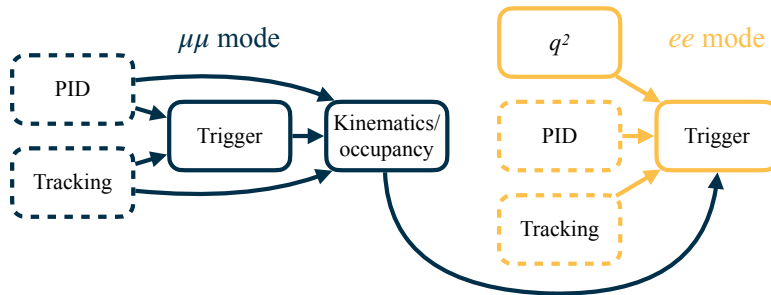


Figure 6.5: Schematic of the order in which corrections to dimuon (blue) and dielectron (yellow) MC samples are computed. Corrections with dashed and solid outlines are obtained from calibration samples and data/MC samples of this measurement, respectively.

The kinematic/occupancy correction must be evaluated using the analysis data and MC samples and is computed in the dimuon mode, as the difference between the reconstructed and true momenta of the B^+ candidates is much larger in the dielectron mode. This requires the corrections of the track reconstruction and PID efficiencies, which are evaluated using calibration samples. The trigger efficiency corrections must be evaluated

Table 6.2: Dielectron and dimuon yields from the fits in each data block (shown in Figs. C.1 & C.1), and their respective ratios.

Block	Quantity	$B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$	Part. reconstructed	Combinatorial
1	ee yield	238230 ± 520	3840 ± 150	29790 ± 200
	$\mu\mu$ yield	752840 ± 890	192110 ± 600	100160 ± 420
	Ratio	0.31644 ± 0.00078	0.01999 ± 0.00078	0.2974 ± 0.0024
5	ee yield	226300 ± 500	4130 ± 150	30490 ± 210
	$\mu\mu$ yield	760180 ± 900	207300 ± 620	101210 ± 430
	Ratio	0.29769 ± 0.00075	0.01992 ± 0.00073	0.3013 ± 0.0024
6	ee yield	189520 ± 460	3380 ± 130	25770 ± 190
	$\mu\mu$ yield	649030 ± 830	176960 ± 580	86370 ± 390
	Ratio	0.29200 ± 0.00080	0.01910 ± 0.00074	0.2984 ± 0.0026
7	ee yield	134700 ± 390	2460 ± 110	18200 ± 160
	$\mu\mu$ yield	460420 ± 700	130630 ± 490	61810 ± 340
	Ratio	0.29257 ± 0.00096	0.01883 ± 0.00085	0.2945 ± 0.0031
8	ee yield	78730 ± 300	1857 ± 93	10750 ± 120
	$\mu\mu$ yield	275740 ± 540	78420 ± 380	36630 ± 260
	Ratio	0.2855 ± 0.0012	0.0237 ± 0.0012	0.2935 ± 0.0039

directly from the analysis samples, with corrections for the PID and tracking efficiencies (and the q^2 resolution and kinematic/occupancy corrections in the dielectron mode) applied to the MC sample.

The corrections of efficiency effects can be expressed as per-event weights, w_i , determined by the ratio between the efficiency as evaluated data and in MC, *i.e.*,

$$w_i = \frac{\varepsilon_{\text{data}}^i}{\varepsilon_{\text{MC}}^i}. \quad (6.5)$$

The application of the q^2 resolution and kinematic/occupancy corrections are described in Section 6.3.2.

Once these terms have been determined, the total efficiency, *i.e.*, the efficiency terms in Eq. 6.1 are computed as

$$\varepsilon_{\text{tot.}} = \varepsilon_{\text{acc.}} \cdot \frac{\sum_{i=1}^{N_{\text{sel.}}} w_{\text{track}}^i \cdot w_{\text{PID}}^i \cdot w_{\text{trig.}}^i \cdot w_{\text{kin.}}^i}{\sum_{j=1}^{N_{\text{acc.}}} w_{\text{kin.}}^j}, \quad (6.6)$$

wherein the fraction is the selection efficiency corrected by the weights w_i . The sum in the numerator is performed over the reconstructed and selected candidates, whilst the sum in the denominator is performed over the generated events within the LHCb acceptance. This section lays out the evaluation of each component in Eq. 6.6, discussing the geometric acceptance efficiencies in Sec 6.3.1, the data-MC correction terms in the dimuon and dielectron modes in Section 6.3.2, respectively, and the combination of these in Section 6.3.3.

6.3.1 Geometric acceptance efficiencies

The MC simulated samples used to evaluate the efficiency to reconstruct and select events contain only the events which lie within the LHCb detector acceptance, *i.e.*, those in which the final-state particles form an angle in the forward direction with the beam axis between $\theta = 10$ mrad ($\eta \sim 5.3$) and $\theta = 400$ mrad ($\eta \sim 1.6$). The rejection of particles outside this range takes place at generator level and is reported for each generated sample [96]. The geometric acceptance efficiencies for each decay mode and block are listed in Table 6.3. Within statistical fluctuations, these are consistent between ee and $\mu\mu$ modes in each block.

Table 6.3: Geometric acceptance efficiencies, $\varepsilon_{\text{acc.}}(\ell\ell)$, of both decay modes in each data-taking block.

Block	$\varepsilon_{\text{acc.}}(ee)$	$\varepsilon_{\text{acc.}}(\mu\mu)$
1	0.17317 ± 0.00052	0.17456 ± 0.00049
5	0.17424 ± 0.00052	0.17449 ± 0.00052
6	0.17334 ± 0.00051	0.17576 ± 0.00052
7	0.17284 ± 0.00056	0.17405 ± 0.00056
8	0.17345 ± 0.00056	0.17272 ± 0.00056

As the ratios of these efficiencies between the dielectron and dimuon modes are very close to unity, this factor is neglected here, with the efficiencies reported hereafter evaluated without this factor.

6.3.2 Data-MC corrections

As described above, the selection efficiencies in each mode must be corrected for differences between the data and MC samples. Corrections are determined accounting for these differences in track reconstruction, PID and trigger efficiencies, and kinematic and occupancy distributions in both modes, and in the q^2 resolution in the dielectron mode.

Track reconstruction efficiency corrections

The tracking efficiencies are determined in data and MC samples from dedicated calibration samples of $B^+ \rightarrow J/\psi(\ell^+\ell^-)K^+$ per the tag-and-probe method applied in Section 4.2.4 in Run 2 and incorporated into the `CalibMon` monitoring framework in Section 5.1.2 in Run 3, wherein the tag lepton is reconstructed as a long track and the probe lepton as a partial track, which is matched to long tracks in the event. The efficiency is extracted through a simultaneous fit to the dilepton mass spectra of probe tracks which are and are not matched to long tracks, performed individually in bins of the probe ℓ^\pm p_T and η . A map of the data-MC corrections is computed as the ratio of the efficiency maps in data and MC. Only data-MC differences in the reconstructing the lepton tracks are considered as, to first order, the data-MC differences in the reconstruction of the kaon track should be identical between the two decay modes and hence cancel in the ratio of Eq. 6.6. For illustration, the correction maps for tracking efficiency data-MC differences in each mode in Block 8 are shown in Fig. 6.6.

The bins of both correction maps are consistently below 1, except for the highest bin in p_T and η , though this is also one of the least well-populated bins. This indicates that the track reconstruction is more efficient in MC, with the corrections providing a

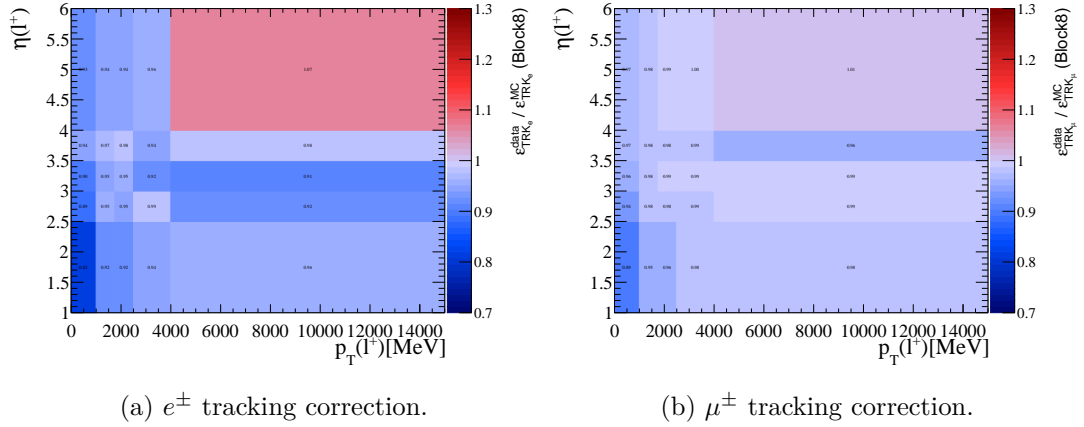


Figure 6.6: Data-MC correction maps for the efficiency to reconstruct each ℓ^\pm as a long track in Block 8.

reducing to account for this overestimation. This affect is greater in the dielectron mode, where the deviations from 1 are greater than those of the dimuon mode; this is shown in Section 6.3.3 to be $\sim 5\%$ per electron and $\sim 2\%$ per muon on average.

The correction maps are applied to each lepton in each event, matching the p_T and η of these leptons to the corresponding bins in the map. The per-event correction factor, w_{track}^i , is then computed as the product of these, *i.e.*, $w_{\text{track}}^i = w_{\text{track}}^{i,\ell^+} \cdot w_{\text{track}}^{i,\ell^-}$.

PID efficiency corrections

Corrections to data-MC differences in the efficiencies of PID requirements are also evaluated from dedicated calibration samples, by the tag-and-probe approaches discussed in Sections 4.2.2 & 5.1.2 and implemented in PIDCa1ib2 [104]. Similarly to the tracking efficiency corrections, these are determined from the ratio of efficiency maps in data and MC; however, as PID requirements are made on all final state particles, these corrections must be evaluated per-particle.

The PID efficiencies of muons are determined with detached $J/\psi \rightarrow \mu^+\mu^-$ decays, those of electrons with $B^+ \rightarrow J/\psi (e^+e^-) K^+$ decays, and those of kaons with $D^{*+} \rightarrow D^0\pi^+$ decays, all of which are evaluated in bins of the respective particle p_T and η . The PID efficiencies of electrons differ significantly depending on whether any recovered bremsstrahlung photons are associated with the electron, hence the electron PID efficiencies are evaluated separately for electrons with and without associated bremsstrahlung. The *sPlot* method [105] is used to obtain the tag-and-probe yields for muons and kaons, whereas for electrons, the strong kinematic dependence of bremsstrahlung effects in the background components requires a fit-and-count approach to be applied, as discussed in Section 5.2.2. As the kaon and muon calibration samples feature higher statistics, a more granular binning scheme can be applied than is applied to the electron samples. The resulting data-MC correction maps of electrons (in both cases), muons and kaons in Block 8 are shown in Fig. 6.7.

In most bins, except those of the kaon correction maps at high p_T/η , the corrections are again below 1, indicating that the PID efficiency is again overestimated in MC. The correction maps of electrons with no associated recovered bremsstrahlung feature the strongest correction factors, particularly for bins at higher η . In contrast, the correction maps of electrons with associated bremsstrahlung are very flat with values lying close

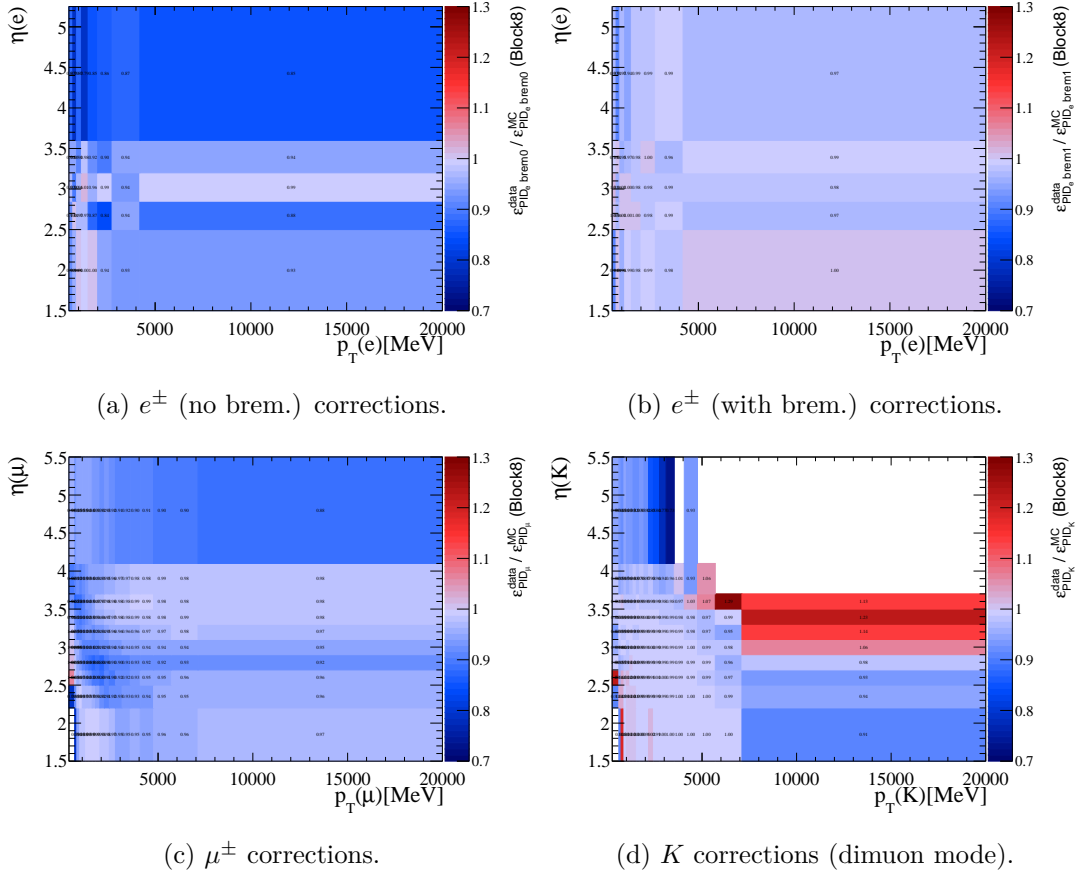


Figure 6.7: PID efficiency data-MC correction maps from Block 8.

to 1. The muon correction maps are relatively flat for harder muons; however, an η -dependent region of worse agreement at low p_T is observed. On average, the K^+ PID correction in both modes is $\sim 1\%$ and cancels in the ratio of Eq. 6.6. The PID corrections of each lepton are $\sim 4\%$ per electron and $\sim 6\%$ per muon, *i.e.*, the dimuon mode is corrected more strongly.

The correction maps are again applied to the MC samples to obtain per-particle per-event correction weights. As for the tracking efficiency corrections, the per-event PID efficiency corrections, w_{PID}^i , are calculated as the product of the per-particle weights in each event, $w_{\text{PID}}^i = w_{\text{PID}}^{i,\ell^+} \cdot w_{\text{PID}}^{i,\ell^-} \cdot w_{\text{PID}}^{i,K^+}$.

Trigger efficiency corrections

Whilst the tracking and PID efficiencies can be determined from calibration samples, calibration samples are not provided for trigger efficiencies (specifically those of HLT1), as these efficiencies are sensitive to the offline selections applied. Instead, these are computed on the analysis data and MC samples. As the trigger response in MC is sensitive to other data-MC differences, the MC samples must be corrected for these differences. The PID and tracking corrections discussed above are applied to both modes; in the dielectron mode, the corrections to the dielectron mass resolution and kinematic/occupancy corrections discussed later are also applied.

The HLT1 efficiencies are evaluated using `TriggerCalib` (described in Section 5.2.2 and in Ref. [164]), applying the TISTOS method to determine the efficiencies in the data and

trigger/PID-correction weighted MC samples in bins of the B^+ p_T and η . Specifically, these are the combined TOS efficiencies of the HLT1TrackMVA and HLT1TwoTrackMVA trigger lines. As for the tracking and PID efficiency correction maps, the trigger efficiency data-MC correction maps are taken as the ratio of these, such as is shown in Fig. 6.8 in Block 8.

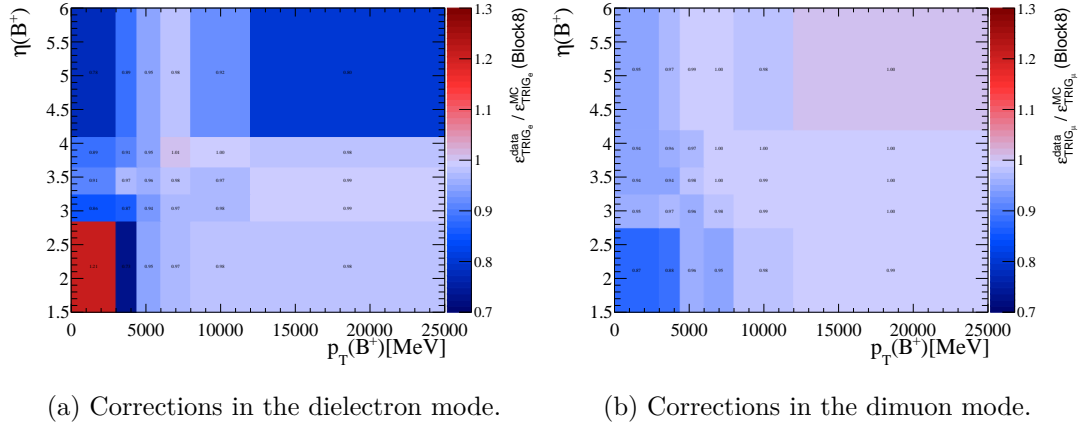


Figure 6.8: Trigger efficiency correction map in Block 8.

The trigger correction maps in the dielectron mode shows significant variation over the B^+ kinematics, with the greatest deviations from 1 arising at low p_T and very low and high η . These are regions in which the trigger efficiency itself varies most rapidly (see Section 5.2.3), though these are also amongst the most sparsely populated bins. For harder B^+ candidates, and those further from the edges of the LHCb acceptance in η , *i.e.*, where the trigger efficiencies are flattest, the corrections are also flattest and lie much closer to 1. Averaging over the MC sample kinematics, corrections of $\sim 2\%$ are obtained in both the dielectron and dimuon mode.

In the dimuon mode, the trigger efficiencies vary slightly less rapidly with p_T and η , resulting in much flatter correction maps. Again, the greatest variations are seen in the lowest bins in p_T and η , where the variation in trigger efficiency is most rapid.

Dielectron mass resolution correction

The quality with which the dielectron invariant mass, or more specifically the resolution of q^2 , is modelled in MC is highly dependent on the precision with which detector effects and bremsstrahlung photons are modelled. Mismodelling of these factors can be accounted for by smearing the q^2 distribution MC to improve the agreement with the resolution measured in the data, incorporating the ratio of resolutions, $f_\sigma = \sigma_{\text{data}}/\sigma_{\text{MC}}$, and difference in positions, $\delta_\mu = \mu_{\text{data}} - \mu_{\text{MC}}$, of the dilepton invariant mass peaks between data and MC. The smeared dilepton mass, m_{smeared} , is then evaluated on a per-event basis as

$$m_{\text{smeared}} = m_{\text{true}} + f_\sigma (m_{\text{reco}} - m_{\text{true}}) + \delta_\mu + (1 - f_\sigma) (\mu_{\text{MC}} - m_{J/\psi}), \quad (6.7)$$

wherein m_{true} and m_{reco} are the generator level (without final state radiation) and reconstructed dilepton masses in simulation, respectively, and $m_{J/\psi}$ is the measured mass of the J/ψ , taken as (3096.900 ± 0.006) MeV/ c^2 from Ref. [5].

To extract the parameters $\sigma_{\text{data,MC}}$ and $\mu_{\text{data,MC}}$, and by extension f_σ and δ_μ , NLL fits to the reconstructed dilepton invariant mass in data and MC are performed. The signal contribution is modelled by a DSCB and the combinatorial background contribution by an exponential. To improve modelling accuracy, the fits, and consequently the smearing, is performed separately for the cases where neither (0 brems.), only one (1 brems.), or both (2 brems.) of the electrons have associated recovered bremsstrahlung. The fits are performed for each data and MC sample and bremsstrahlung case in each block, with these fits and their comparisons shown in Fig. 6.9. The resulting values of δ_μ , f_σ and μ_{MC} are listed for all blocks in Table C.3 in Appendix C.2.1.

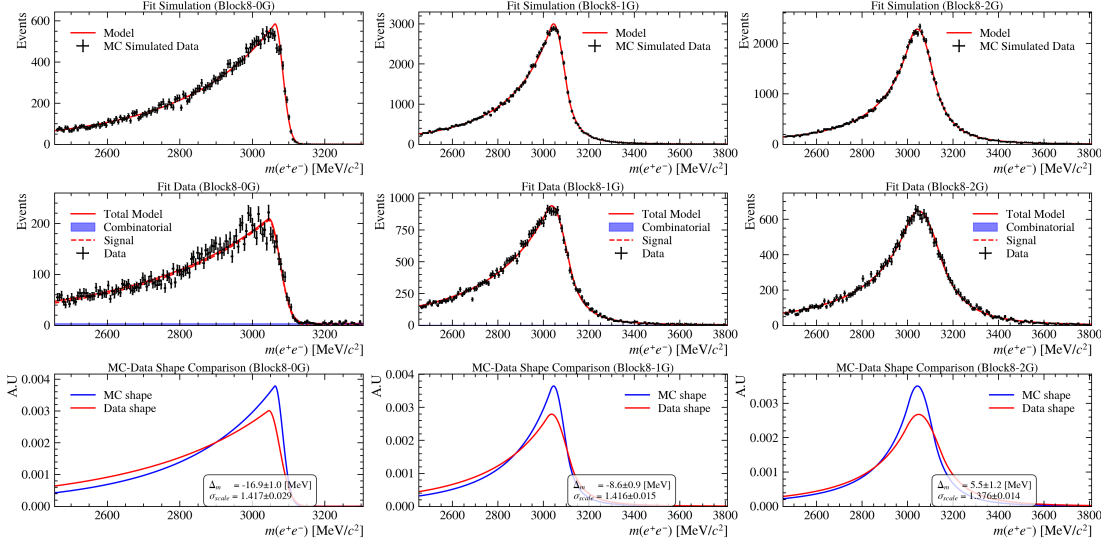


Figure 6.9: Fits to the reconstructed dilepton invariant mass in each bremsstrahlung case in Block 8 MC (top row) and data (middle row) samples, and their comparisons (bottom row).

To validate that the smearing has the intended effect, the reconstructed and smeared dilepton invariant mass spectra in the MC samples are compared to the *sWeighted* spectrum in data, as shown in Fig. 6.10 for Block 8. Application of the smearing results in a close agreement between data and MC, with the 0 brems. case in most agreement and the greatest disagreement seen in the 1 brems. case. The resolutions of the distributions and right-hand tails appear to agree well with the data; however, the far-left tail is still subject to small data-MC disagreements.

The q^2 correction is incorporated into Eq. 6.6 by applying the q^2 window described in Section 6.1.3 to the smeared q^2 distribution, which modifies $N_{\text{sel.}}$ in the dielectron mode. This modified the efficiencies in the dielectron mode by $\sim 8\%$; as no equivalent correction is applied to the dimuon mode, this correction has the greatest impact on the ratio of efficiencies and, by extension, $r_{J/\psi}$.

Kinematic and occupancy corrections

Finally, the mismodelling of kinematic properties of the final-state particles and detector occupancy in simulation is corrected by training a BDT classifier (the `GBRweighter` algorithm from `hep_ml` [119]) to divide the feature space into regions of similar data-MC agreement. In this measurement, these features are taken to be the p , p_T and η of the B^+ candidate, and the number of PVs and long tracks present in the event. The score of

6 Measurement of $r_{J/\psi}$ in 2024 LHCb data

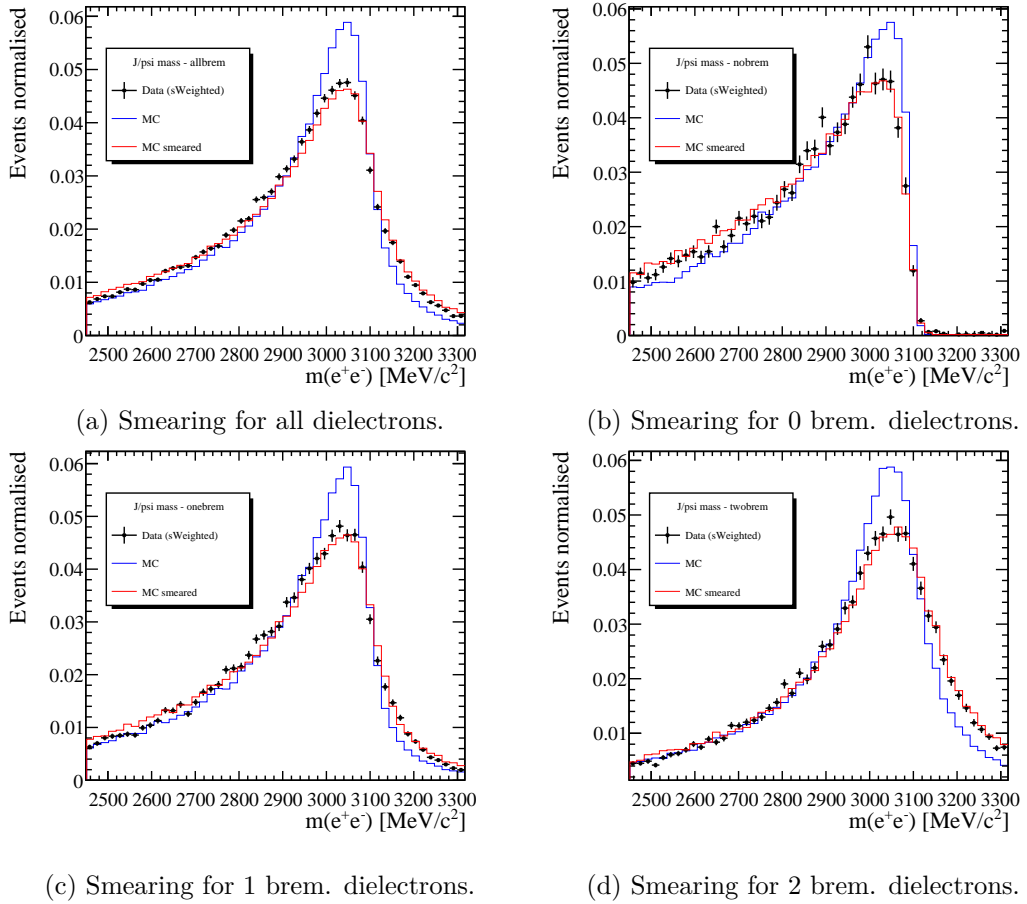


Figure 6.10: Comparison of the reconstructed and smeared dielectron invariant mass spectra in MC with the reconstructed dielectron mass spectrum in data, for Block 8.

a trained classifier on a given MC event is itself the per-event weight, w_{kin}^i , necessary to correct for the data-MC differences; applying these weights to MC samples brings them into closer agreement with the data. As the dimuon mode has higher statistics and significantly narrower J/ψ and B^+ mass resolutions, the BDT is trained on the dimuon data and MC samples, and applied to both modes. To factor out data-MC differences in tracking, PID and trigger efficiencies, the respective correction factors (as derived above) are applied to the MC samples included in the training.

The application of a BDT trained on the dimuon mode in Block 8 to the dimuon and dielectron MC samples in Block 8 is shown in Fig. 6.11, binned in one-dimensional projections in p_T and η of the B^+ candidate. The BDT brings the MC in both modes into very good agreement with the majority of deviations between the corrected MC and data appearing to arise in statistical fluctuations. The effect of this correction is $\sim 1\%$ in the dimuon mode and $\sim 2\%$ in the dielectron mode.

6.3.3 Combination of efficiencies and corrections

The data-MC corrections determined in Sections 6.3.1 & 6.3.2, respectively, are combined in each mode according to Eq. 6.6 to give the corrected selection efficiency (as per Section 6.3.1, $\varepsilon_{\text{acc.}}$ is neglected). The statistical uncertainties on the efficiencies must account for the correlations between the corrections. A bootstrapping approach to the

6.3 Estimation of efficiencies and data-MC corrections

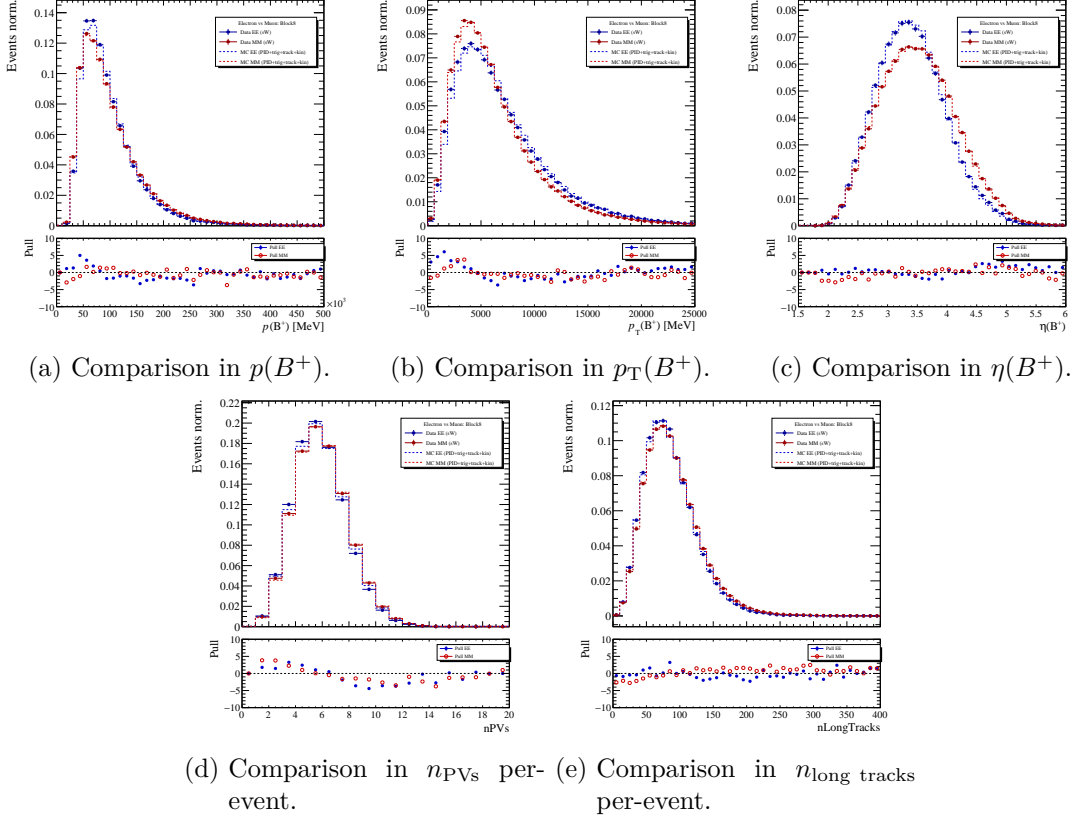


Figure 6.11: Application of the kinematic/occupancy correction BDT trained on dimuon Block 8 data and MC (corrected for data-MC differences in tracking, PID and trigger efficiencies), to Block 8 dimuon and dielectron samples.

statistical uncertainties is taken, wherein each event is assigned 100 weights, w_{BS} , are drawn at random from a Poisson distribution of mean 1 using a common seed dependent only on the run number and event number associated with each event [178]. These can be incorporated for the reconstructed/selected and generated e MC events, such that Eq. 6.6 is modified to

$$\varepsilon_{tot.}^k = \varepsilon_{acc.} \cdot \frac{\sum_{i=1}^{N_{sel.}} w_{BS}^{i,k} w_{track}^i \cdot w_{PID}^i \cdot w_{trig.}^i \cdot w_{kin.}^i}{\sum_{j=1}^{N_{acc.}} w_{BS}^{j,k} w_{kin.}^j}. \quad (6.8)$$

If the approach is stable, the resulting distributions of values of $\varepsilon_{tot.}^k$ should follow Gaussian distributions. If this is the case, then the mean and standard deviation of these resulting values can be taken as the central values and statistical uncertainties of each efficiency, respectively.

The bootstrapping is performed across each data block and mode, with the resulting relative uncertainties (defined as $s_\varepsilon^k = (\varepsilon^k - \bar{\varepsilon}) / \bar{\varepsilon}$) in Block 8 shown in Fig. 6.12. Both distributions are Gaussian, with no clear skew to either side; the distribution in the dimuon mode is notably narrower than that of the dielectron mode, owing to the higher statistics in the dimuon mode in both calibration samples and the samples of this measurement.

6 Measurement of $r_{J/\psi}$ in 2024 LHCb data

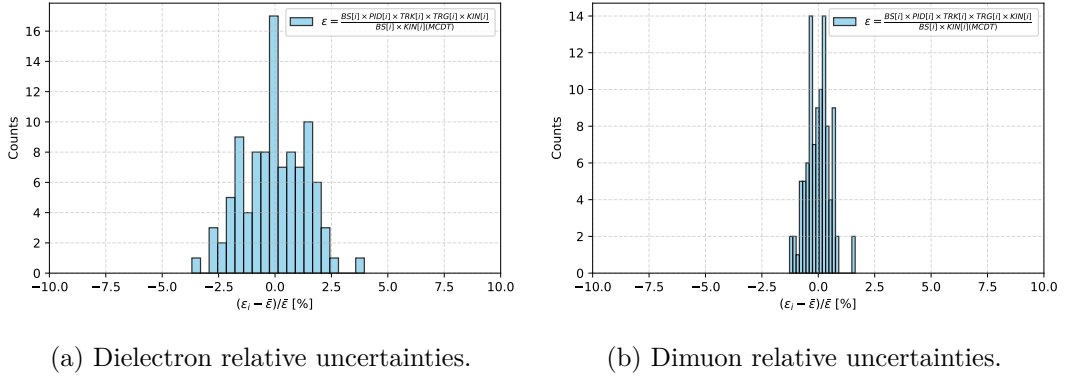


Figure 6.12: Distributions of relative uncertainties from bootstrapping approach applied to the Block 8 efficiencies in each mode.

The effect of applying each correction factor one-by-one is studied by starting from an efficiency in which all correction weights are set to 1, then allowing each to take its respective value and determining the efficiency in each case. Performing precisely this study, the efficiencies in the dielectron and dimuon modes, and their ratio, are shown in Fig. 6.13. All statistical uncertainties evaluated in the study are computed via the bootstrapping method described above.

In both modes, the application of each correction reduces the efficiency; this is unsurprising, as the correction maps shown above consistently demonstrate that each contribution to the efficiency is overestimated in MC. Applying the corrections together, the efficiencies in the dimuon and dielectron modes are reduced by $\sim 19\%$ and $\sim 29\%$, respectively. This results in an average modification of the ratio of efficiencies by $\sim 10\%$. The equivalent modification of the efficiency ratio in Run 2 was $\sim 20 - 30\%$ [77, 78]; where the greater initial difference in the dimuon and dielectron modes was driven by large data-MC differences in the L0 trigger response.

6.4 Consideration of systematic uncertainties

The systematic uncertainties on the constituents of $r_{J/\psi}$ have not yet been evaluated; however, this section lays out the systematic effects relevant to this measurement and how these will be evaluated.

The primary systematic effect present in the determination of the yields arises from the choice of model for each component. To estimate the scale of this effect, pseudoexperiments can be generated replacing the model of each component one-by-one, and fits to the generated datasets performed with both the nominal model and alternative model. The relative difference between the two yields can then be taken as the corresponding systematic uncertainty. Replacing the signal model evaluates effects in the resolution description, particularly in the radiative tail of the dielectron signal. Replacing the background models evaluates whether these more empirical models accurately describe their corresponding components.

In the efficiencies, systematic uncertainties arise primarily from the limited sizes of the samples and the choice of binning scheme used to determine the correction maps. To estimate the scale of this for each correction map, 100 maps are generated from the original map, in which each bin is varied randomly within its statistical uncertainty. The efficiencies according to Eq. 6.6 for each of these varied maps, and the efficiency




Figure 6.13: Efficiencies in the dielectron and dimuon modes, and their ratio, for each data block and for the successive inclusion of each correction factor.

ratio taken; the relative difference between the mean of these ratios and the equivalent ratio evaluated in Section 6.3.2 is taken as the corresponding systematic uncertainty. The corrections can be treated as statistically independent, such that the total systematic uncertainty on the efficiency ratios are taken as the quadrature sum of the systematic uncertainties from each correction.

Further contributions to each correction should also be considered. In the PID corrections, the choice of binning scheme (granularity and choice of variables) and correlations with event occupancy not described by the calibration samples may also be significant. For the tracking corrections, significant other contributions include differences in momentum resolution causing migration of events between bins, imperfect modelling of interactions with the detector material, multiple scattering, and energy loss, and kinematic dependence not fully described by the calibration samples. The trigger corrections, in addition to the effects described in Section 5.2.2, may be dependent on event occupancy. The choice of input features and reweighter used to determine the kinematic/occupancy corrections should be studied. Finally, for the q^2 smearing, the modelling of the long radiative tail in the dilepton invariant mass, the difference between cuts to the B^+ mass with and without the smearing applied, *e.g.*, in HLT2, and statistical fluctuations in the parameters of Eq. 6.7 may all have an impact.

6.5 Summary of results and conclusion

Bringing together the yields from Section 6.2 and efficiencies/corrections from Section 6.3, $r_{J/\psi}$ is computed (per Eq. 6.1) for each block of data. The values of $r_{J/\psi}$ are shown, incorporating each correction factor from Section 6.3.2, with corresponding statistical uncertainties in Fig. 6.14. The total correction of $r_{J/\psi}$, which is $\sim 10\%$, is much smaller than the equivalent correction in Run 2, as the removal of the L0 trigger significantly reduces the difference in the kinematics of the dielectron and dimuon modes.




Figure 6.14: Values of $r_{J/\psi}$ evaluated incorporating each correction from Section 6.3.2. The equivalent is shown for the dielectron and dimuon efficiencies, and their ratio, in Fig. 6.13.

All results presented above are stated with their statistical uncertainties; the systematic uncertainties are still to be evaluated but will follow in due course. As previously observed in Run 2, the systematic uncertainties associated with the measurement of a single ratio can extend up to a scale of $\sim 5\%$. With such a scale of systematic uncertainty on $r_{J/\psi}$ (as is to be expected here), the values of $r_{J/\psi}$ reported in Fig. 6.14 would all be consistent with 1. This systematic uncertainty is then reduced in a double ratio measurement, and could be expected to lie at or below $\sim 1\%$.

In its current status, this measurement already demonstrates that the application of analysis techniques, especially those for determining corrections to the efficiencies and decay properties, to the 2024 data can be performed successfully.

7 Conclusion

This thesis follows the evolution of the LHCb trigger with two measurements. The world-first measurement of $\langle A_{\text{untagged}}^s \rangle$ was performed, discussed in Chapter 4, studying the pp collision data collected by the LHCb experiment in 2016-18 to determine, in each D_s^- decay mode,

$$\langle A_{\text{untagged}}^s \rangle_{KK\pi\pi} = (-1.2 \pm 6.2 \pm 1.0) \times 10^{-3},$$

$$\langle A_{\text{untagged}}^s \rangle_{\pi\pi\pi\pi} = (-0.4 \pm 1.6 \pm 0.2) \times 10^{-2}.$$

These results agree extremely well with the SM expectation, *i.e.*, consistency with 0, deviating only by 0.19σ and 0.25σ in the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively. This sheds new light on the $b \rightarrow c\bar{u}q$ anomaly, as the strong agreement of these results with the SM can be used to further constrain possible NP contributions, which had been anticipated to induce CP asymmetries up to the scale of $\mathcal{O}(10^{-2})$. The measurement is statistically limited, with the statistical sensitivity suffering most notably from the L0 trigger efficiency ($\varepsilon \sim 40\%$) and the HCAL separation cut ($\varepsilon \sim 25\%$, necessitated by L0). Removal of the hardware (L0) trigger, and hence both of these contributions, could improve the statistical sensitivity by $\sim 3\times$; incorporating also the increase in integrated luminosity between Run 2 and Run 3, an equivalent measurement in Run 3 could potentially expect improvements of the statistical sensitivity of up to $\sim 6\times$.

The upgrade of the LHCb detector, discussed in Chapter 5, saw, as a highlight, the implementation of an entirely software-based trigger. Selection algorithms such as those discussed were successfully deployed in the upgraded triggers, with the widespread use of real-time monitoring ensuring that necessary retuning was performed as soon as possible. The data-driven TISTOS method, implemented in `TriggerCalib`, was used to evaluate the efficiencies of HLT1 and the HLT1-HLT2 combination in data taken in 2024 and 2025. The efficiencies of HLT1 in 2024 and 2025 are compared with their equivalents in Run 2, *i.e.*, including the L0 trigger efficiency, demonstrating conclusively that the improvements in trigger yields on which the trigger upgrade was motivated were realised. In hadronic channels, global trigger efficiency improvements of up to $\sim 3\times$ were observed, with local trigger efficiencies at low p_T improved by even greater factors. The efficiencies of HLT1 and HLT2 combined demonstrate that the b -topo. trigger provides a dilution of these gains for events saved to the full stream; however, this dilution is significantly reduced by the successful retuning of the b -topo. trigger in early 2025.

The measurement of $r_{J/\psi}$ in $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^+$ decays in 2024 pp collision data, described in Chapter 6, demonstrates the direct impact of the improvements in the detector and trigger on physics analyses. In particular, corrections to the selection efficiencies of each mode which account for data-MC differences are significantly smaller than their Run 2 equivalents. This is primarily driven by the removal of the L0 trigger, which previously introduced large differences in the kinematics of the two modes. The measured values of $r_{J/\psi}$ in each block, which are dominated by systematic uncertainties, are all expected to be consistent with 1.

In the course of this thesis, the necessity of the removal of the L0 trigger was demonstrated through the measurement of $\langle A_{\text{untagged}}^s \rangle$ in 2016-18 data (as described in Chapter 4), the sensitivity of which suffers from the low L0 trigger efficiency for hadronic decay

7 Conclusion

modes. The removal of the L0 trigger and the redevelopment of the trigger selection algorithms to perform efficient selections at much higher event rates was successful, with the anticipated improvements in trigger efficiencies realised. Finally, the measurement of $r_{J/\psi}$ in 2024 data demonstrates that the data collected by the upgraded LHCb detector and trigger is of a high quality, with its features understood to the degree necessary to perform precision measurements in Run 3 of the LHC.

Perhaps I'm old and tired, but I think that the chances of finding out what's actually going on are so absurdly remote that the only thing to do is to say, "Hang the sense of it," and keep yourself busy. I'd much rather be happy than right any day.

"The Hitchhiker's Guide to the Galaxy", *Douglas Adams*

A Appendix for the Run 2 measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$

This appendix serves to provide additional background information on the measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$ decays laid out in Chap. 4. The samples of pp collision data studied in this measurement are listed by data-taking year and magnet polarity in Tab. A.1.

Table A.1: LHCb pp collision data samples studied in the measurement of $\langle A_{\text{untagged}}^s \rangle$ described in Chap. 4.

Year	Magnet polarity	Int. lumi. / fb^{-1}
2016	Up	0.81
	Down	0.86
2017	Up	0.83
	Down	0.88
2018	Up	1.14
	Down	1.05
Total	Up	2.79
	Down	2.79

A.1 Updated results

Since the submission of this thesis, the measurement described in Chap. 4 has been completed. The corresponding paper has been submitted to *Phys. Rev. Lett.*, and a preprint has been made publicly available [179]. The results presented in Ref. [179] supersede those presented in this thesis.

A.2 Per-event weighting of MC samples

As described in Section 4, the per-event PID and L0 trigger efficiencies are attached to each event and applied as per-event weights. The effect of applying these weights (individually and together) upon the $hh\pi\pi$ and $hh\pi$ invariant masses, and the p_T and η of the associated π^+ and D_s^- are shown in Figs. A.1 & A.2 for the $KK\pi\pi$ and $\pi\pi\pi\pi$ final states, respectively.

A Appendix for the Run 2 measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$

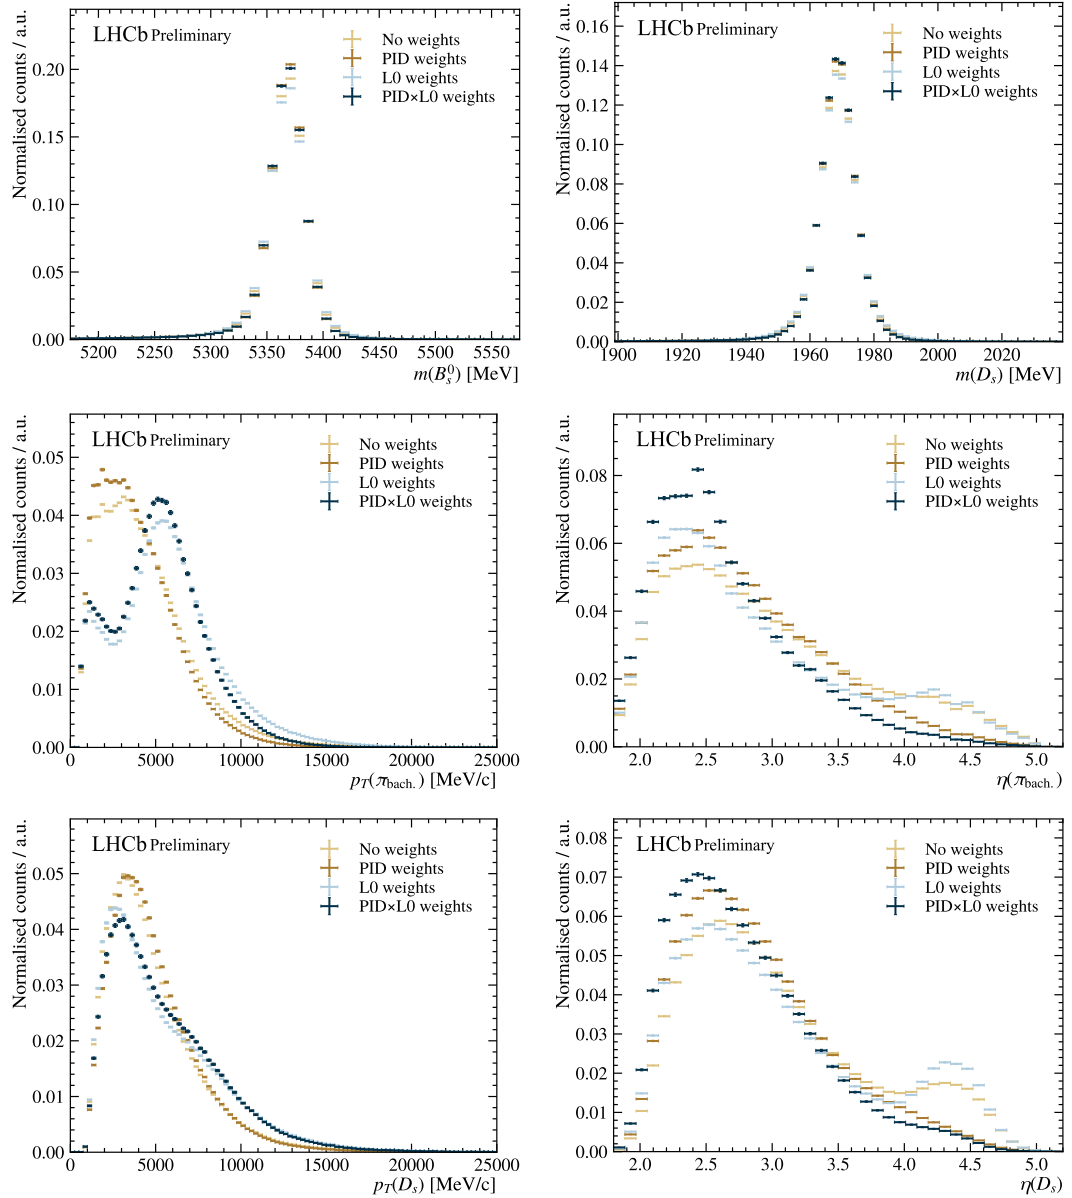


Figure A.1: Impact of per-event weights from PIDCalib2 and L0HadronTables on key distributions in the $KK\pi\pi$ final state.

A.2 Per-event weighting of MC samples

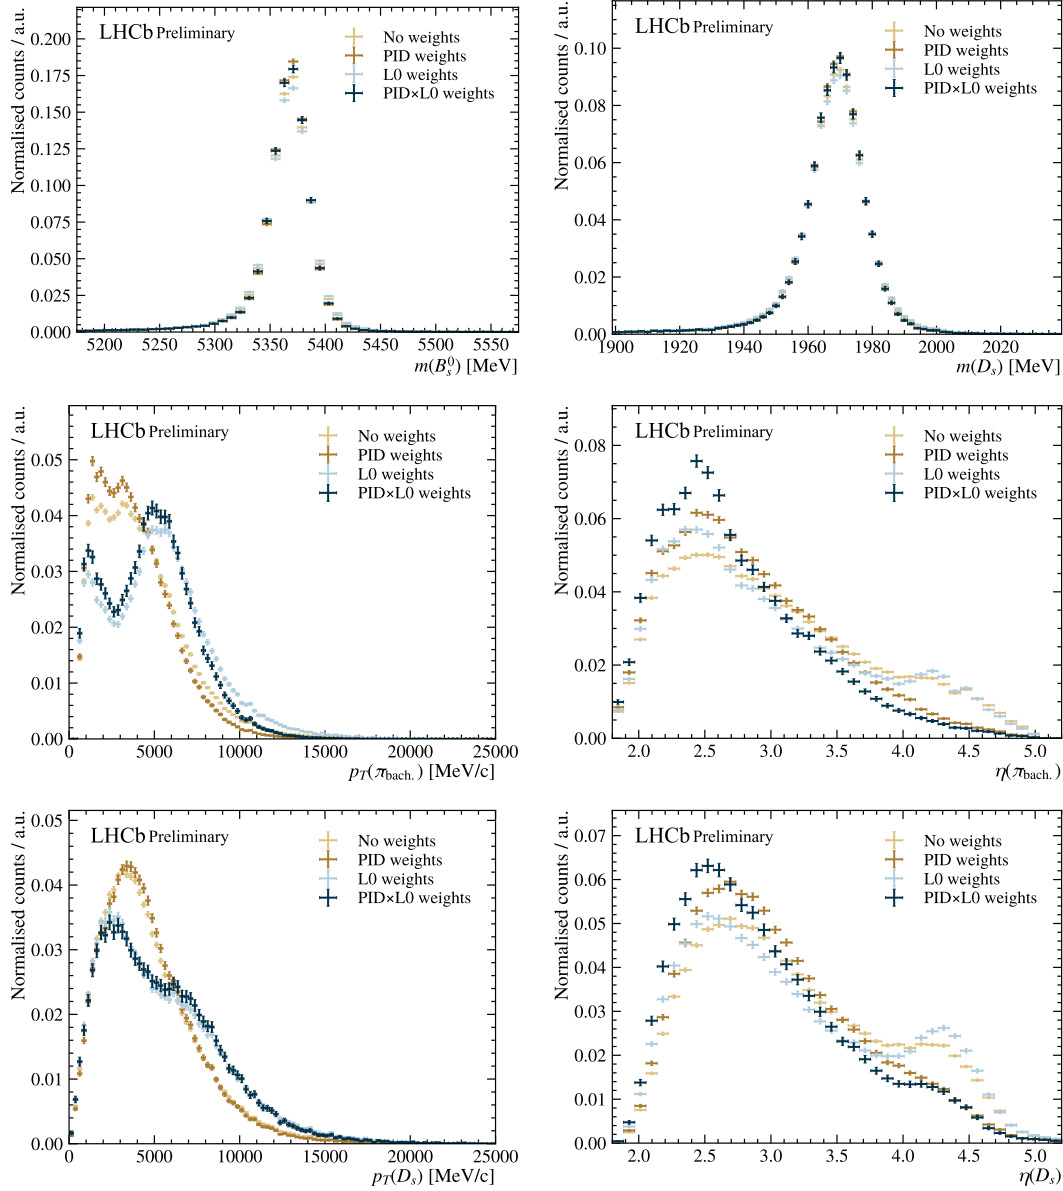


Figure A.2: Impact of per-event weights from PIDCalib2 and L0HadronTables on key distributions in the $\pi\pi\pi\pi$ final state.

A.3 Parameterisation and results of the invariant mass fits

The invariant mass fits performed to extract A_{raw} for each final state are described in Section 4.3. This section provides additional information on the specific signal and combinatorial background models applied/trialled, and provides an additional visualisation of the nominal invariant mass fits.

A.3.1 Signal models

A wide range of statistical distributions are commonly used to model invariant mass peaks in studies of b -hadron decays. This section serves as a whistle-stop tour of the distributions discussed in this thesis, and particularly in Chap. 4.

The Crystal Ball function, named for the Crystal Ball experiment, first introduced in Ref. [180], consists of a piecewise combination of a Gaussian core with a power-law tail. The double-sided Crystal Ball function is a modification of this such that the function is subject to a piecewise division, with differing cores and tails on either side of the central value of the cores, μ :

$$f_{\text{DSCB}}(m) = \begin{cases} A_L \cdot (B_L - \frac{m-\mu}{\sigma_L})^{-n_L}, & \text{for } \frac{m-\mu}{\sigma_L} < -\alpha_L \\ \exp\left(-\frac{1}{2} \cdot \left[\frac{m-\mu}{\sigma_L}\right]^2\right), & \text{for } \frac{m-\mu}{\sigma_L} \leq 0 \\ \exp\left(-\frac{1}{2} \cdot \left[\frac{m-\mu}{\sigma_R}\right]^2\right), & \text{for } \frac{m-\mu}{\sigma_R} \leq \alpha_R \\ A_R \cdot (B_R + \frac{m-\mu}{\sigma_R})^{-n_R}, & \text{otherwise,} \end{cases} \quad (\text{A.1})$$

where σ_i are widths of each Gaussian core, n_i determine power-law behaviour of the tails, and A_i , B_i normalise the function as

$$A_i = \left(\frac{n_i}{|\alpha_i|}\right)^{n_i} \cdot \exp\left(-\frac{|\alpha_i|^2}{2}\right), \quad (\text{A.2a})$$

$$B_i = \frac{n_i}{|\alpha_i|} - |\alpha_i|. \quad (\text{A.2b})$$

The Hypatia function is Crystal Ball-like (, a Gaussian-like core with power law tails), but instead uses a hyperbolic core [181]. The two-sided Hypatia function is defined similarly to the double-sided Crystal Ball function, as

$$f_{\text{Hy}}(m) = \begin{cases} G(\mu - a_1\sigma) / \left(1 - \frac{m}{n_1 G(m)/G'(\mu - a_1\sigma)}\right)^{n_1} & \text{if } \frac{m-\mu}{\sigma} < -a_1 \\ \left((m-\mu)^2 + A_\lambda^2(\zeta)\sigma^2 \right)^{\frac{1}{2}\lambda - \frac{1}{4}} e^{\beta(m-\mu)} K_{\lambda-\frac{1}{2}} \left(\zeta \sqrt{1 + \left(\frac{m-\mu}{A_\lambda(\zeta)\sigma}\right)^2} \right) \equiv G(m) & \text{otherwise} \\ G(\mu + a_2\sigma) / \left(1 - \frac{m}{-n_2 G(m)/G'(\mu + a_2\sigma)}\right)^{n_2} & \text{if } \frac{m-\mu}{\sigma} > a_2 \end{cases}, \quad (\text{A.3})$$

where μ and σ are the mean and width of the core (G), a_i and n_i describe the tails, and λ , ζ , β parameterise the core shape.

The extended Cruijff function is laid out in Ref. [182] and is defined as a Gaussian of modified width

$$f_{\text{EC}}(m|\mu, \beta, \sigma_L, \sigma_R, \alpha_L, \alpha_R) = \begin{cases} C_L(m|\mu, \beta, \sigma_L, \alpha_L) & m \leq \mu \\ C_R(m|\mu, \beta, \sigma_R, \alpha_R) & m > \mu \end{cases}, \quad (\text{A.4})$$

A.3 Parameterisation and results of the invariant mass fits

where $C_{L,R}(m|\mu_{L,R}, \sigma_{L,R}, \alpha_{L,R}, \beta)$ are given by

$$C_{L,R}(m|\mu_{L,R}, \sigma_{L,R}, \alpha_{L,R}, \beta) = A \exp\left(\frac{-(m-\mu)^2(1+\beta(m-\mu)^2)}{2\sigma_{L,R}^2 + \alpha_{L,R}(x-\mu)^2}\right), \quad (\text{A.5})$$

which reduces to a regular Cruijff function in the case $\beta = 0$.

The Johnson S_U function, $f_{\text{Jo}}(m)$ is defined from the transformation of a normal distribution as

$$m \rightarrow m' = \gamma + \delta N\left(\frac{m-\mu}{\lambda}\right), \quad (\text{A.6})$$

such that $f_{\text{Jo}}(m)$ is given by

$$f_{\text{Jo}}(m) = \frac{\delta}{\lambda\sqrt{2\pi}} \frac{1}{\sqrt{1+\left(\frac{m-\mu}{\lambda}\right)^2}} \exp\left[-\frac{1}{2}\left(\gamma + \delta \sinh^{-1}\left(\frac{m-\mu}{\lambda}\right)\right)^2\right], \quad (\text{A.7})$$

wherein μ and λ are analogous to the mean and width of a Gaussian distribution [183], and δ and γ further specify the transformation.

These distributions can themselves be sufficient to model invariant mass peaks, though often these are combined to provide descriptions with more degrees of freedom, such as the sum of a Hypatia and Johnson modelling $B_s^0 \rightarrow D_s^- \pi^+$ decays in Section 4.3 and the sum of a DSCB and Gaussian modelling $B^+ \rightarrow J/\psi(\mu^+ \mu^-) K^+$ decays in Section 6.2.

Table A.2: Parameters of Hypatia-Johnson PDFs fit to $B_s^0 \rightarrow D_s^- \pi^+$ MC samples, per Figs. 4.7a & 4.7b.

Parameter	Fit value	
	$KK\pi\pi$ final state	$\pi\pi\pi\pi$ final state
Hypatia μ / MeV	5369.440 ± 0.045	5372.7 ± 1.2
Hypatia σ / MeV	62.1 ± 2.1	41 ± 11
Hypatia α_L	0.2434 ± 0.0074	0.0995 ± 0.049
Hypatia n_L	2.53 ± 0.31	2.34 ± 0.71
Hypatia α_R	0.4015 ± 0.0060	0.42 ± 0.11
Hypatia n_R	3.038 ± 0.061	1.89 ± 0.53
Hypatia λ	-8.70 ± 0.14	-6.8 ± 1.7
Hypatia-Johnson c_{HyJo}	0.315 ± 0.025	0.102 ± 0.040
Johnson δ	3.271 ± 0.011	2.11 ± 0.23
Johnson γ	0.0862 ± 0.0039	0.286 ± 0.078
Johnson λ	45.84 ± 0.19	33.8 ± 3.9

A Appendix for the Run 2 measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$

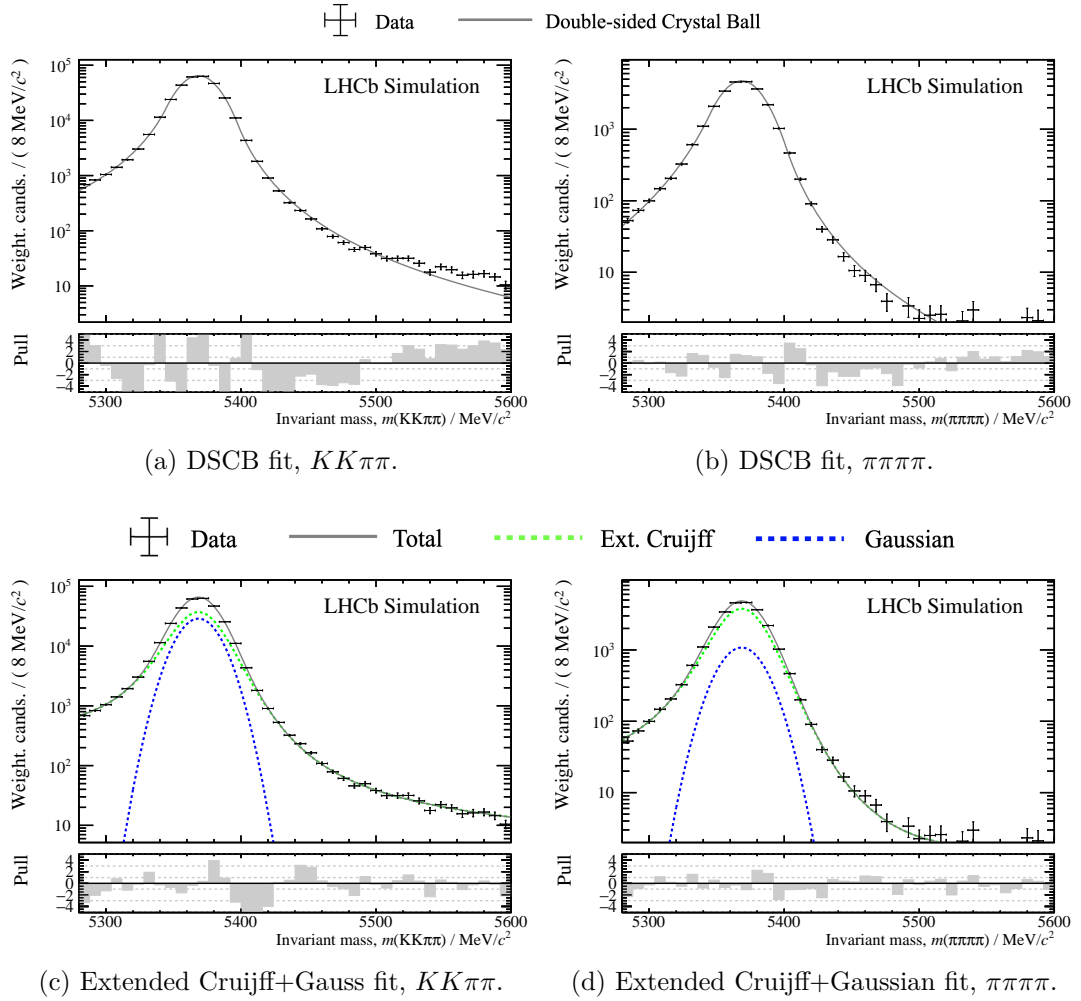


Figure A.3: Fits of the DSCB and extended Cruijff+Gaussian models to the 2016-18 combined $B_s^0 \rightarrow D_s^- \pi^+$ MC samples in both $hh\pi\pi$ final states. Equivalent fits using the Hypatia+Johnson S_U model are shown in Fig. 4.7 in Section 4.3.1.

A.3.2 Combinatorial models

The single exponential combinatorial background model described in Section 4.3.1 is also fit to the combinatorial background data samples, as shown in Fig. A.4.

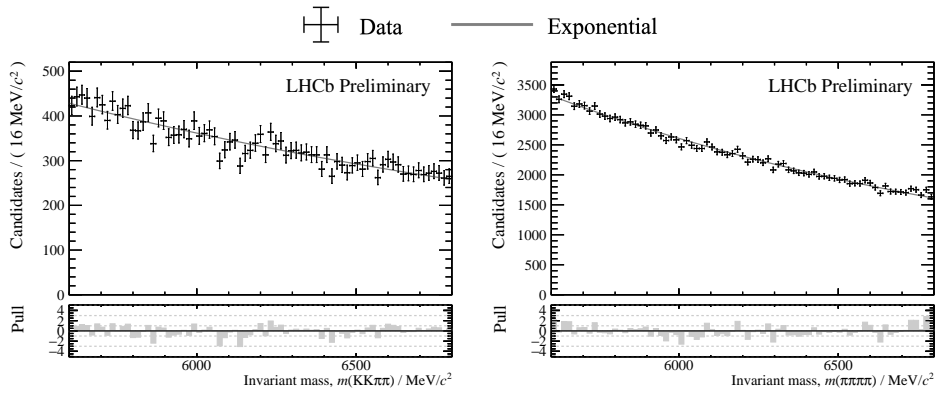


Figure A.4: Fits of single exponential PDF to the combinatorial background data samples (2016-18, both magnet polarities). The equivalent fits using exponential+constant models are shown in Fig. 4.10 in Section 4.3.1.

A.3.3 Results of the nominal invariant mass fit

The nominal fit performed in Section 4.3.2 is shown in Fig. 4.11 on a logarithmic y -axis scale. The same fits are shown here in Fig. A.5 on linear y -axis scales.

A Appendix for the Run 2 measurement of $\langle A_{\text{untagged}}^s \rangle$ in $B_s^0 \rightarrow D_s^- \pi^+$

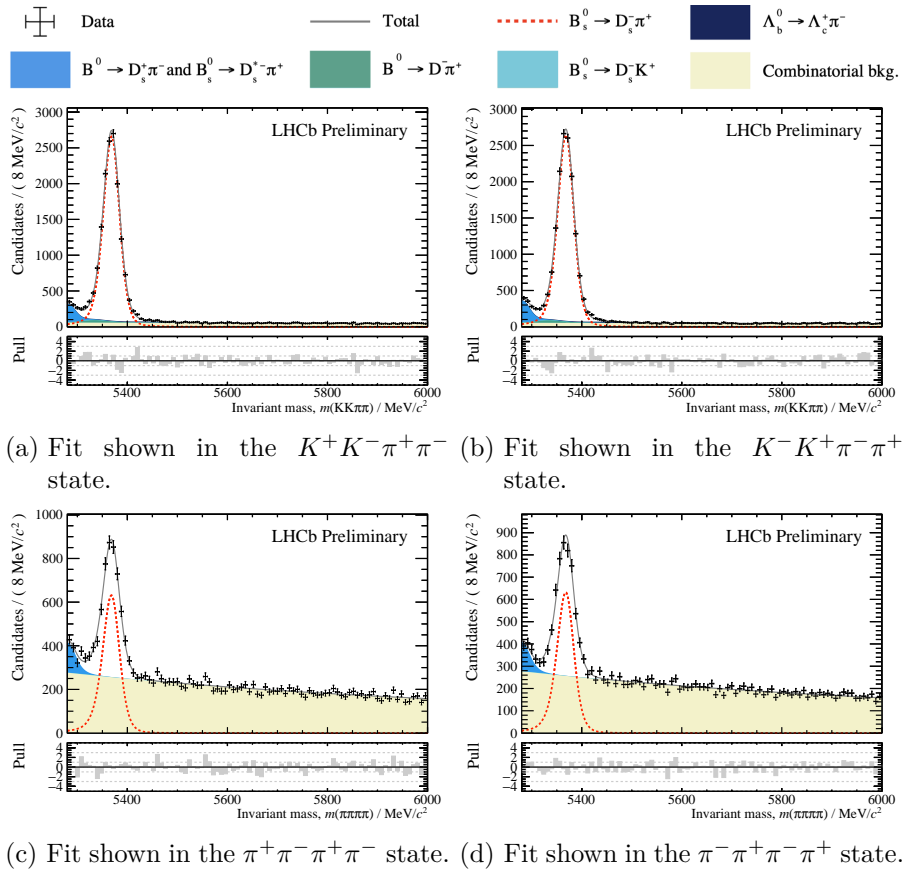


Figure A.5: Nominal invariant mass fit to each final state, shown for each charge conjugate state. Corresponding parameters are listed in Tab. 4.7 and the same fits shown on a logarithmic scale in Fig. 4.11, both of which are in Section 4.3.2.

B Appendix to the LHCb trigger in Run 3

This appendix provides additional information on Chap. 5, which describes the LHCb detector and trigger upgrades, and the data-driven TISTOS method for computing trigger efficiencies, before then evaluating the efficiencies of HLT1 and HLT2 against the equivalent Run 2 efficiencies. The data samples taken in 2024 and 2025, including those used in evaluating the efficiencies of the upgraded trigger are listed in Tab. B.1. Note that data-taking blocks have not been assigned in 2025.

Table B.1: Blocks of stable data-taking configurations in 2024 and 2025 (up to the June technical stop), listed in chronological order. Preliminary numbers taken from Ref. [184].

Year	Block	μ	Polarity	Int. luminosity / fb^{-1}
2024	4	≤ 4.4	Down	1.39
	3		0.90	
	2		0.61	
	1	Up	1.18	
	5		1.16	
	6	Down	1.00	
	7		0.74	
	8	Up	0.44	
2025	—	5.3	Down	0.47
	—		Up	1.74

B.1 Selections of the cut-based inclusive detached dilepton trigger

In Section 5.2.1, the $CB\ell\ell$ trigger was developed ahead of 2023 data-taking and retuned for 2024 data-taking to improve on the performance observed in 2023. The cut-based selection requirements applied before and after this retuning are listed in Tabs. B.2 & B.3, respectively.

Table B.2: Selection cuts of the CB $l\ell$ trigger as deployed in 2023 data-taking. Note that for two-body candidates, the b candidate selections are used.

Candidate	Selection cut
Leptons	$p_T > 0 \text{ MeV}/c$
	$p > 0 \text{ GeV}/c$
	$\chi_{\text{IP}}^2 > 36$
Electrons	$\text{PID}_e > 2$
Muons	$\text{PID}_\mu > 2$
Dilepton candidate	$\chi_{\text{vtx.}}^2/N_{\text{dof}} < 9$
	$\chi_{\text{DOCA}}^2 < 30$
	$\chi_{\text{FD}}^2 > 300$
Additional track	$p_T > 250 \text{ MeV}/c$
	$p > 2 \text{ GeV}/c$
	$\chi_{\text{IP}}^2 > 9$
b candidate	$p_T > 0 \text{ MeV}/c$
	$\chi_{\text{vtx.}}^2/N_{\text{dof}} < 9$
	$\chi_{\text{IP}}^2 < 16$
	$\chi_{\text{FD}}^2 > 300$

B.1 Selections of the cut-based inclusive detached dilepton trigger

Table B.3: Selection cuts of the retuned CB $\ell\ell$ trigger, as deployed in data-taking from 2024 onwards. Note that for two-body candidates, the dilepton candidate selections are used.

Candidate	Selection cut
Leptons	$p_T > 400 \text{ MeV}/c$
	$p > 2 \text{ GeV}/c$
	$\chi_{\text{IP}}^2 > 36$
Electrons	$\text{PID}_e > 2$
	ECAL acceptance
Muons	$\text{PID}_\mu > 0$
	isMuon
Dilepton candidate	$p_T > 400 \text{ MeV}/c$
	$p > 2 \text{ GeV}/c$
	$m(\ell\ell) \in [0, 6.5] \text{ GeV}/c^2$
	$\chi_{\text{FD}}^2 > 64$
	$\chi_{\text{vtx.}}^2/N_{\text{dof}} < 20$
Dielectron/electron-muon	$\chi_{\text{DOCA}}^2 < 10$
Dimuon	$\chi_{\text{DOCA}}^2 < 30$
Additional track	$p_T > 400 \text{ MeV}/c$
	$p > 2 \text{ GeV}/c$
	$\chi_{\text{IP}}^2 > 36$
b candidate	$p_T > 0 \text{ MeV}/c$
	$\chi_{\text{vtx.}}^2/N_{\text{dof}} < 16$
	$\chi_{\text{FD}}^2 > 64$

B.2 Retuning of the topological b -trigger

The retuning of the b -topo. trigger is also reported in Section 5.2.1. The features on which the b -topo. NNs are defined, including the monotonicity requirements before and after retuning, are listed in Tab. B.4. The preselection which b -topo. candidates are subject to prior to the application of the NNs, which were not changed during the retuning, are listed in Tab. B.5

Table B.4: Input features of the two- and three-body Lipschitz NNs of the b -topo. trigger. Variables which are monotonic only in 2024 (both in 2024 and 2025) are denoted by * (\dagger).

Two-body Lipschitz NN features	Three-body Lipschitz NN features
	Min. p_T of 2-body cand. children \dagger
	Sum of p_T of 2-body cand. children*
	p_T of 2-body cand.
	M_{corr} of 2-body cand.
	$\chi_{\text{end vtx.}}^2/N_{\text{dof}}$ of 2-body cand.
	Max. χ_{FD}^2 of 2-body cand.
	Max. $\text{DOCA}_{\text{end vtx.}}$ of 2-body cand.
Min. χ_{IP}^2 of 2-body cand. children \dagger	Min. χ_{IP}^2 of 3-body cand. children \dagger
Max. χ_{IP}^2 of 2-body cand. children	Max. χ_{IP}^2 of 3-body cand. children
	χ_{IP}^2 of 2-body cand.
	Min. p_T of 3-body cand. children \dagger
	Sum of p_T of 3-body cand. children*
	p_T of 3-body cand.*
	M_{corr} of 3-body cand.
	$\chi_{\text{end vtx.}}^2/N_{\text{dof}}$ of 3-body cand.
	Max. χ_{FD}^2 of 3-body cand.
	Max. $\text{DOCA}_{\text{end vtx.}}$ of 3-body cand.

Table B.5: Preselection applied to two- and three-body candidates.

Target	Selection requirement
Final state particle	$p_T > 250 \text{ MeV}/c$
	$p > 3 \text{ GeV}/c$
	$\chi_{\text{IP}}^2 > 4$
Two-body candidate	DOCA $< 0.2 \text{ mm}$
	$p_T > 1 \text{ GeV}/c$
	$\chi_{\text{FD}}^2 > 20$
	$\max\{\chi_{\text{IP}}^2\}_{\text{tracks}} > 20$
	$\chi_{\text{vtx.}}^2/N_{\text{dof}} < 10$
	DIRA > 0.9
	$M_{\text{corr}} \in [1, 12] \text{ GeV}/c^2$
$\eta \in [2, 5]$	
Any track hasisMuon (for muonic 2-body line)	
Three-body candidate	DOCA $< 0.3 \text{ mm}$
	$p_T > 2 \text{ GeV}/c$
	$\chi_{\text{FD}}^2 > 20$
	DIRA > 0.9
	$M_{\text{corr}} \in [1, 12] \text{ GeV}/c^2$
$\eta \in [2, 5]$	
Any track hasisMuon (for muonic 3-body line)	

C Appendix to the 2024 measurement of $r_{J/\psi}$

This final appendix related to the measurement of $r_{J/\psi}$ in 2024 pp collision data described in Chap. 6.

C.1 Requirements of exclusive HLT2 Turbo lines

Table C.1: Cut-based selection requirements of the HLT2 line H1t2RD_BuToKpEE for the blocks studied in Chap. 6.


Object	Variable	Selection requirement	
		Block 1	Blocks 5-8
B^+ candidate	$m_{\text{inv.}}$	$\in [4.5, 7.0] \text{ GeV}/c^2$	$\in [3.5, 7.0] \text{ GeV}/c^2$
	FD χ^2		> 100
	Vertex χ^2/ndof		< 4
	IP χ^2		< 25
	DIRA		> 0.9999
	p_{T}	$> 0 \text{ MeV}/c$	$> 500 \text{ MeV}/c$
Isolation	ΔR^2		< 0.25
Dielectrons	DOCA χ^2		< 36
	FD χ^2		> 36
	Vertex χ^2/ndof		< 9
	p_{T}		$> 0 \text{ MeV}/c$
Electrons	IP χ^2		> 25
	p_{T}	$> 350 \text{ MeV}/c$	$> 500 \text{ MeV}/c$
	p		$> 0 \text{ MeV}/c$
	PID _e		> 0.0
	Upstream electrons allowed		True
	Combination policy	—	≥ 1 long (Block 7-8)
Kaons	IP χ^2		> 9
	p_{T}	$> 400 \text{ MeV}/c$	
	PID _k		> -4

C.2 Extraction of yields


In Section 6.2, the invariant mass fits from which the yields in each mode are obtained are only shown for Block 8. The fits to the blocks 1, 5, 6 and 7 are shown in Fig. C.1.

Table C.2: Cut-based selection requirements of H1t2RD_BuToHpMuMu_Incl_Bp for the blocks studied in Chap. 6.


Object	Variable	Selection requirement Blocks 1, 5-8
B^+ candidate	$m_{\text{inv.}}$	$\in [4.5, 7.1] \text{ GeV}/c^2$
	p_T	$> 0 \text{ MeV}/c$
	FD χ^2	> 16
	Vertex χ^2/ndof	< 36
	IP χ^2	< 36
	DIRA	> 0.9995
Dimuons	$m_{\text{inv.}}$	$\in [0, 7.1] \text{ GeV}/c^2$
	DOCA χ^2	< 36
	FD χ^2	> 9
	Vertex χ^2/ndof	< 9
Muons	p_T	$> 0 \text{ MeV}/c$
	IP χ^2	> 3
	p_T	$> 300 \text{ MeV}/c$
	p	$> 0 \text{ MeV}/c$
	PID $_{\mu}$	> -4
Hadrons	IP χ^2	> 3
	p_T	$> 300 \text{ MeV}/c$
	p	$> 0 \text{ MeV}/c$
	PID $_k$	No cut




(a) Fit to dielectron mode, Block 1




(b) Fit to dimuon mode, Block 1




(c) Fit to dielectron mode, Block 5




(d) Fit to dimuon mode, Block 5




(e) Fit to dielectron mode, Block 6



(f) Fit to dimuon mode, Block 6



(g) Fit to dielectron mode, Block 7



(h) Fit to dimuon mode, Block 7

 Figure C.1: Fits to the B^+ DTF mass in samples for data blocks 1 and 5-7. Corresponding yields are listed in Tab. 6.2. Equivalent fits in Block 8 are shown in Fig. 6.4.

C.2.1 Dielectron mass resolution smearing

The smearing of the dielectron mass resolution, as described in Section 6.3.2, requires fits to the dielectron invariant mass in data and MC samples for each data block and bremsstrahlung case. Whilst Fig. 6.9 shows these fits for Block 8, the full set of smearing parameters for each year and bremsstrahlung case are listed in Tab. C.3.

Table C.3: Relative positions, δ_μ , and widths, f_σ , and MC positions of invariant mass peaks, μ_{MC} , from fits to the reconstructed dilepton invariant mass in data and MC samples in each bremsstrahlung case.

Block	Brem. case	$\delta_\mu / \text{MeV}/c^2$	f_σ	$\mu_{\text{MC}} / \text{MeV}/c^2$
1	0 brem.	-15.91 ± 0.72	1.403 ± 0.022	3062.3 ± 1.21
	1 brem.	-12.36 ± 0.49	1.393 ± 0.008	3047.0 ± 1.38
	2 brem.	-8.06 ± 0.64	1.390 ± 0.008	3041.3 ± 1.05
5	0 brem.	-14.73 ± 0.61	1.359 ± 0.016	3062.5 ± 0.33
	1 brem.	-13.82 ± 0.52	1.362 ± 0.009	3043.2 ± 1.34
	2 brem.	-7.48 ± 0.71	1.378 ± 0.009	3040.9 ± 1.11
6	0 brem.	-15.93 ± 0.64	1.373 ± 0.018	3067.1 ± 1.32
	1 brem.	-11.37 ± 0.56	1.370 ± 0.010	3046.0 ± 1.30
	2 brem.	-5.25 ± 0.71	1.371 ± 0.009	3040.4 ± 1.03
7	0 brem.	-16.78 ± 0.78	1.448 ± 0.023	3066.6 ± 1.41
	1 brem.	-10.13 ± 0.70	1.389 ± 0.011	3040.8 ± 1.35
	2 brem.	4.71 ± 0.95	1.349 ± 0.011	3041.9 ± 1.22
8	0 brem.	-16.90 ± 1.05	1.417 ± 0.029	3063.0 ± 0.35
	1 brem.	-8.60 ± 0.90	1.416 ± 0.015	3045.2 ± 1.22
	2 brem.	5.51 ± 1.15	1.376 ± 0.014	3043.6 ± 1.03

List of Tables

4.1	Summary of background contributions which may enter the $hh\pi\pi$ invariant mass distributions, and the sources by which they arise. Contributions marked * and † are proven not present in (see Section 4.1.4) and removed entirely (see Section 4.1.3) from the data samples, respectively.	38
4.2	Summary of selection requirements.	40
4.3	Efficiencies of kinematic selections, calculated with and without distinction on final-state charges.	48
4.4	Efficiencies of PID requirements, calculated with and without distinction on final-state charges.	49
4.5	Efficiencies of L0Hadron, calculated with and without distinction on final-state charges.	50
4.6	Fractions describing the yields of $B_s^0 \rightarrow D_s^- K^+$ and $B^0 \rightarrow D^- \pi^+$ relative to the yield of $B_s^0 \rightarrow D_s^- \pi^+$ (per Eq. 4.15) and its constituent components. Kinematic and PID efficiencies are taken from MC corresponding to 2018 magnet up data-taking.	57
4.7	Parameters of the nominal invariant mass fits shown in Fig. 4.11.	59
4.8	Parameters resulting from the charge-agnostic fit, as shown in Fig. 4.15	61
4.9	Charge-specific selection efficiencies and corresponding asymmetries from $B_s^0 \rightarrow D_s^- \pi^+$ MC samples.	68
4.10	Charge-specific PID asymmetries in $B_s^0 \rightarrow D_s^- \pi^+$, in the (non-)resonant $D_s^+ \rightarrow h^+ h^- \pi^+$ phase space regions, and their luminosity-weighted means.	69
4.11	Charge-specific L0 selection efficiencies and corresponding asymmetries from $B_s^0 \rightarrow D_s^- \pi^+$ MC samples.	69
4.12	Tracking detection asymmetry in $KK\pi\pi$ and its components.	72
4.13	Tracking detection asymmetry in $\pi\pi\pi\pi$ and its components.	73
4.14	Summary of systematic uncertainties on A_{raw}	75
4.15	Summary of systematic uncertainties on A_{det}	77
4.16	Differences between the nominal values of $\langle A_{\text{untagged}}^s \rangle$ and values determined in subsamples of the datasets.	78
4.17	Summary of the components of $\langle A_{\text{untagged}}^s \rangle$, for the combination of all years and magnet polarities.	78
5.1	Per channel weights applied in the minimisation of Eq. 5.14 for the optimisation of b -topo. trigger NN thresholds in 2025.	101
6.1	Summary of selection requirements.	120
6.2	Dielectron and dimuon yields from the fits in each data block (shown in Figs. C.1 & C.1), and their respective ratios.	124
6.3	Geometric acceptance efficiencies, $\varepsilon_{\text{acc.}}(\ell\ell)$, of both decay modes in each data-taking block.	125
A.1	LHCb pp collision data samples studied in the measurement of $\langle A_{\text{untagged}}^s \rangle$ described in Chap. 4.	139
A.2	Parameters of Hypatia-Johnson PDFs fit to $B_s^0 \rightarrow D_s^- \pi^+$ MC samples, per Figs. 4.7a & 4.7b.	143

List of Tables

B.1	Blocks of stable data-taking configurations in 2024 and 2025 (up to the June technical stop), listed in chronological order.	147
B.2	Selection cuts of the $CB\ell\ell$ trigger as deployed in 2023 data-taking. Note that for two-body candidates, the b candidate selections are used.	148
B.3	Selection cuts of the retuned $CB\ell\ell$ trigger, as deployed in data-taking from 2024 onwards. Note that for two-body candidates, the dilepton candidate selections are used.	149
B.4	Input features of the two- and three-body Lipschitz NNs of the b -topo. trigger. Variables which are monotonic only in 2024 (both in 2024 and 2025) are denoted by * (\dagger).	150
B.5	Preselection applied to two- and three-body candidates.	151
C.1	Cut-based selection requirements of the HLT2 line $H1t2RD_BuToKpEE$ for the blocks studied in Chap. 6.	153
C.2	Cut-based selection requirements of $H1t2RD_BuToHpMuMu_Incl_Bp$ for the blocks studied in Chap. 6.	154
C.3	Relative positions, δ_μ , and widths, f_σ , and MC positions of invariant mass peaks, μ_{MC} , from fits to the reconstructed dilepton invariant mass in data and MC samples in each bremsstrahlung case.	156

List of Figures

2.1	The fundamental particles of the SM, with electric charges (upper left) and spins (upper right) annotated.	4
2.2	CKM unitarity triangle defined by the angles in Eqs. 2.27a-2.27c, with the constraints imposed by relevant measurements overlaid.	11
2.3	Schematic of the topology of a b -hadron decaying to a c -hadron. Distances not drawn to scale.	12
2.4	Mixing of B_s^0 and \bar{B}_s^0 mesons as measured by LHCb in $B_s^0 \rightarrow D_s^- \pi^+$ decays in Run 2.	13
2.5	Leading order Feynman diagrams for $b \rightarrow s \ell^+ \ell^-$ transitions in the SM. . .	15
2.6	Overview of experimental (blue) and theoretical (yellow) averages of branching fractions of $B_{(s)}^0 \rightarrow D_{(s)}^{(*)-} h^+$ decays, courtesy of Nicole Skidmore. . .	19
2.7	Summaries of the current state of $b \rightarrow s \ell^+ \ell^-$ anomalies.	20
3.1	Schematic diagram of the LHC and accompanying accelerator facilities. . .	21
3.2	$b\bar{b}$ production cross-section dependence on the angle/pseudorapidity of each quark.	22
3.3	Side-on schematic of the LHCb detector in Run 2 of the LHC.	23
3.4	Performance of long track (see Fig. 3.5) reconstruction in Run 2.	24
3.5	Top-down schematic of the track types defined in LHCb. The upstream tracker (UT) and scintillating fibre (SciFi) tracker, employed in Run 3, are introduced in Section 5.1.1. Muon tracks, not shown in this schematic, are simply long tracks with hits in M1-M5.	24
3.6	Track and vertex reconstruction performance, including comparisons between datasets from Run 2 and the 2012 dataset.	25
3.7	Cherenkov angle in RICH2 as a function of track momentum.	26
3.8	Efficiencies and fake rates of PID_k and ProbNNk variables at loose and strict working points, as evaluated during Run 2.	28
3.9	Schematic of the LHCb Run 2 dataflow.	28
3.10	Schematic of the LHCb Run 2 trigger systems.	29
3.11	Efficiencies of the L0 trigger selections on key channels of interest to the LHCb physics programme.	30
3.12	Efficiencies of HLT1TrackMVA and HLT1TwoTrackMVA on key channels of interest to the LHCb physics programme.	31
3.13	Alignment and calibration procedure for a typical LHC fill in Run 2, indicating the points at which sufficient data has been accumulated for each process.	32
3.14	Efficiencies of the topological b -hadron trigger in HLT2 on key channels of interest to the LHCb physics programme.	32
3.15	Schematic of the Turbo event model, in which objects can be persisted as needed, ranging from only the HLT2 candidate (top) to the full event (bottom).	33

List of Figures

4.1	Distributions of $hh\pi\pi$ and $hh\pi$ invariant masses, and Dalitz plots of $hh\pi$ phase space in the 2016-18 dataset, following the application of the selection requirements listed in Table 4.2. The $KK\pi\pi$ ($\pi\pi\pi\pi$) final state is shown on the left (right). Windows around the ϕ and K^{*0} in the $D_s^- \rightarrow K^-K^+\pi^-$ mode and $f_0(980)$ in the $D_s^- \rightarrow \pi^-\pi^+\pi^-$ mode described in Section 4.1.1 are overlaid as dashed lines.	41
4.2	$hh\pi\pi$ and $hh\pi$ invariant mass distributions of $B^0 \rightarrow D^-\pi^+$, $\Lambda_b^0 \rightarrow \bar{\Lambda}_c^-\pi^+$ and $B_s^0 \rightarrow D_s^-\rho^+$, with the $hh\pi\pi$ and $hh\pi$ mass windows overlaid as applicable.	43
4.3	$hh\pi\pi$ invariant masses of candidates with $m(hh\pi) > 2010$ MeV, subject to varying requirements on D_s^+ flight distance.	45
4.4	Invariant masses of the combination of the associated π^+ with each pair of hadrons from the D_s^- decay, for the $KK\pi\pi$ (left) and $\pi\pi\pi\pi$ (right) final states.	46
4.5	HCAL separation of each pair of final-state hadrons.	47
4.6	Effects of per-event weights from <code>PIDCalib2</code> and <code>LOHadronTables</code> , and their combination, on key distributions in the $KK\pi\pi$ final state.	52
4.7	Fits of the Hypatia+Johnson S_U model to the 2016-18 combined $B_s^0 \rightarrow D_s^-\pi^+$ MC samples in both $hh\pi\pi$ final states. Equivalent fits using the DSCB and extended Cruijff+Gaussian models are shown in Fig. A.3 in Appendix A. The corresponding parameters are listed in Table A.2.	53
4.8	Comparison of peaking background components occupying similar spaces in the $KK\pi\pi$ invariant mass.	54
4.9	Template PDFs modelling each peaking background contribution.	55
4.10	Fits of the exponential+constant model to the combinatorial background data samples (2016-18, both magnet polarities). The equivalent fits of a single exponential model are shown in Fig. A.4 in Appendix A.	56
4.11	Nominal invariant mass fit to each final state, shown for each charge conjugate state. Corresponding parameters are listed in Table 4.7. The same fits are shown on a linear scale in Fig. A.5 in Appendix A.	58
4.12	Correlations between parameters of the nominal fits, as listed in Table 4.7 and shown in Fig. 4.11.	59
4.13	Pulls and pseudoexperiment errors on A_{raw} from the nominal pseudoexperiment study. Gaussian distributions fit to each of these are overlaid.	61
4.14	Widths and means of Gaussian distributions as fit to the pulls of each floating parameter in the nominal fit.	62
4.15	Charge-agnostic equivalent to the nominal invariant mass fits shown in Fig. 4.11. Corresponding parameters are listed in Table 4.8.	63
4.16	Charge-agnostic fit to $m(hh\pi\pi)$ using <code>B2DXFitters</code> with the lower limit of the mass window extended to $5100 \text{ MeV}/c^2$	64
4.17	Gaussian means and widths, as fit to $A_{\text{raw}}(KK\pi\pi)$ for scenarios of assigned asymmetries in each background component, with those of the nominal fit, per Fig. 4.13, overlaid (grey/black).	65
4.18	Gaussian means and widths, as fit to $A_{\text{raw}}(\pi\pi\pi\pi)$ for scenarios of assigned asymmetries in each background component, with those of the nominal fit, per Fig. 4.13, overlaid (grey/black).	66
4.19	Weighting of the $D^+ \rightarrow K^-\pi^+\pi^+$ calibration sample to the <i>sWeighted</i> data sample from this analysis, for data taken in 2016 with the magnet up configuration.	72

4.20	Cubic spline function modelling the B_s^0 decay time acceptance, with knots and coefficients.	74
4.21	Dependence of the $A_{\text{det.}}^{\text{ProbNN}}$ on the choice of requirements ProbNNk and ProbNNpi , parameterised by ε . Values of ε corresponding to the PID requirements applied within and outside the resonant D_s^- phase-space are marked by hollow and solid circles, respectively. The uncertainties on $A_{\text{det.}}^{\text{ProbNN}}$ are solely statistical, and are entirely correlated with one another as they study $A_{\text{det.}}^{\text{ProbNN}}$ within the same sample.	76
5.1	Side-on schematic of the LHCb detector in Run 3 of the LHC.	80
5.2	Resolutions of reconstructed PV positions in 2018 and 2024	80
5.3	Performance of K^\pm identification in early 2024 data taking.	81
5.4	Signal yields of the LHCb trigger from Run 2 at instantaneous luminosity working points, normalised to the yields at the Run 1 instantaneous luminosity of $\mathcal{L}_{\text{inst}} = 2 \times 10^{32} \text{ cm}^{-2} \text{ s}^{-1}$	82
5.5	Schematic diagram of the Run 3 LHCb trigger configuration	82
5.6	Schematic diagram of the flow of data through the LHCb trigger in Run 3. The onward flow of data in offline processing is shown in more detail in Fig. 5.10.	83
5.7	Comparison of Allen throughputs for GPUs considered for HLT1.	83
5.8	Performance of the Looking Forward track reconstruction algorithm on simulated samples.	84
5.9	Monitoring web page for the CalibMon monitoring of $K_S^0 \rightarrow \pi^+\pi^-$ for a single LHC fill in early 2024 data-taking.	85
5.10	Dataflow of the LHCb offline processing.	87
5.11	Cumulative integrated luminosity of pp collisions recorded by LHCb, shown by year of data-taking.	88
5.12	Rates and efficiencies in $B_s^0 \rightarrow e^+e^-$ MC of the 2-body dielectron CB$\ell\ell$ line for selection working points in PID$_e$ (each curve, stricter for lighter curves) and electron χ_{IP}^2 (each point per curve). Chosen working point is shown by dashed lines.	93
5.13	Monitoring histograms from HLT2 processing of Runs 270352-270563, comparing the events selected by the MV$\ell\ell$ (light blue) and CB$\ell\ell$ (dark blue) triggers.	94
5.14	Rates and $B^+ \rightarrow J/\psi (\ell^+\ell^-) K^+$ efficiencies of the CB$\ell\ell$ trigger for each random set of topological selection cuts.	95
5.15	Combined trigger efficiencies of the 2023 (original) and 2024 (retuned) implementations of the CB$\ell\ell$ trigger, and the MV$\ell\ell$ trigger as deployed in 2023, computed on signal MC for 7 different final states containing a e^+e^- or $\mu^+\mu^-$ pair.	96
5.16	MVA response across a toy space in features of the b -topo. trigger.	97
5.17	Efficiencies of the b -topo. trigger in late 2024 data, evaluated with the TISTOS method (described later in Section 5.2.2).	98
5.18	Histograms in the space of minimum χ_{FD}^2 of final state tracks and two-body candidate lifetime in the background and signal training samples.	99
5.19	Efficiencies of the b -topo. trigger with optimised NN thresholds for rate working points $R \in [6, 18]$ kHz. Statistical uncertainties shown are fully correlated.	102
5.20	Efficiencies of the 2024 and 2025 b -topo. trigger configurations computed on 2024 MC samples.	103

List of Figures

5.21	Comparison of $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ 2024 MC samples with no selection applied and with a TIS requirement.	105
5.22	Comparison of $\varepsilon_{\text{Trig}}$. (per Eq. 5.19) with true efficiency in $B^+ \rightarrow J/\psi (\mu^+ \mu^-) K^+$ Run 2 MC samples over increasingly granular phase space binnings.	105
5.23	Comparison of background mitigation methods	107
5.24	HLT1(Two)TrackMVA efficiencies for muonic, electronic and hadronic b -hadron decays in early 2024 data, shown in red. Equivalent efficiencies in Run 2 (the product of respective L0 trigger efficiencies with those of HLT1) are shown in blue. The distribution of b -hadron candidates produced in pp collisions which lie within the detector acceptance, taken from MC, is overlaid in grey.	111
5.25	HLT1(Two)TrackMVA efficiencies for b -hadron decays in late 2024 data (red). As in Fig. 5.24, the equivalent Run 2 L0 and HLT1 efficiencies (black) are overlaid.	112
5.26	HLT1(Two)TrackMVA efficiencies for b -hadron decays in early 2025 data. As in Fig. 5.24, the equivalent Run 2 L0 and HLT1 efficiencies and p_T distribution of b -hadrons in the LHCb acceptance are overlaid.	113
5.27	Combined HLT1(Two)TrackMVA and b -topo. trigger efficiencies for b -hadron decays in late 2024 data Equivalent efficiencies from Run 2 (the product of those of the L0, HLT1, and HLT2 triggers) are shown in black.	114
5.28	Combined HLT1(Two)TrackMVA and b -topo. trigger efficiencies for b -hadron decays in 2025 data.	115
6.1	Fits of signal template PDFs to respective $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^+$ MC in each decay mode.	121
6.2	Fits of RooKeysPdf to partially reconstructed $B^+ \rightarrow J/\psi (\ell^+ \ell^-) K^{*0}$ MC samples in each decay mode.	122
6.3	Fits of exponential distributions to the upper $m(B^+)$ sideband in data samples in each decay mode.	122
6.4	Fits to the B^+ DTF mass in Block 8 data samples. Corresponding yields are listed in Table 6.2. Equivalent fits in data blocks 1 and 5-7 are shown in Fig. C.1.	123
6.5	Schematic of the order in which corrections to dimuon (blue) and dielectron (yellow) MC samples are computed. Corrections with dashed and solid outlines are obtained from calibration samples and data/MC samples of this measurement, respectively.	123
6.6	Data-MC correction maps for the efficiency to reconstruct each ℓ^\pm as a long track in Block 8.	126
6.7	PID efficiency data-MC correction maps from Block 8.	127
6.8	Trigger efficiency correction map in Block 8.	128
6.9	Fits to the reconstructed dilepton invariant mass in each bremsstrahlung case in Block 8 MC (top row) and data (middle row) samples, and their comparisons (bottom row).	129
6.10	Comparison of the reconstructed and smeared dielectron invariant mass spectra in MC with the reconstructed dielectron mass spectrum in data, for Block 8.	130
6.11	Application of the kinematic/occupancy correction BDT trained on dimuon Block 8 data and MC (corrected for data-MC differences in tracking, PID and trigger efficiencies), to Block 8 dimuon and dielectron samples.	131

6.12	Distributions of relative uncertainties from bootstrapping approach applied to the Block 8 efficiencies in each mode.	132
6.13	Efficiencies in the dielectron and dimuon modes, and their ratio, for each data block and for the successive inclusion of each correction factor.	133
6.14	Values of $r_{J/\psi}$ evaluated incorporating each correction from Section 6.3.2. The equivalent is shown for the dielectron and dimuon efficiencies, and their ratio, in Fig. 6.13.	134
A.1	Impact of per-event weights from PIDCalib2 and L0HadronTables on key distributions in the $KK\pi\pi$ final state.	140
A.2	Impact of per-event weights from PIDCalib2 and L0HadronTables on key distributions in the $\pi\pi\pi\pi$ final state.	141
A.3	Fits of the DSCB and extended Cruijff+Gaussian models to the 2016-18 combined $B_s^0 \rightarrow D_s^- \pi^+$ MC samples in both $hh\pi\pi$ final states. Equivalent fits using the Hypatia+Johnson S_U model are shown in Fig. 4.7 in Section 4.3.1.	144
A.4	Fits of single exponential PDF to the combinatorial background data samples (2016-18, both magnet polarities). The equivalent fits using exponential+constant models are shown in Fig. 4.10 in Section 4.3.1.	145
A.5	Nominal invariant mass fit to each final state, shown for each charge conjugate state. Corresponding parameters are listed in Tab. 4.7 and the same fits shown on a logarithmic scale in Fig. 4.11, both of which are in Section 4.3.2.	146
C.1	Fits to the B^+ DTF mass in samples for data blocks 1 and 5-7. Corresponding yields are listed in Tab. 6.2. Equivalent fits in Block 8 are shown in Fig. 6.4.	155

Bibliography

- [1] S. Weinberg, *A Model of Leptons*, Phys. Rev. Lett. **19** (1967) 1264.
- [2] A. Salam and J. C. Ward, *Weak and electromagnetic interactions*, Il Nuovo Cimento (1955-1965) **11** (1959) 568.
- [3] S. L. Glashow, *The renormalizability of vector meson interactions*, Nuclear Physics **10** (1959) 107.
- [4] F. Capozzi *et al.*, *Neutrino masses and mixing: Entering the era of subpercent precision*, Phys. Rev. D **111** (2025) 093006, [arXiv:2503.07752](#).
- [5] Particle Data Group, S. Navas *et al.*, *Review of particle physics*, Phys. Rev. D **110** (2024) 030001.
- [6] ATLAS collaboration, G. Aad *et al.*, *Observation of a new particle in the search for the Standard Model Higgs boson with the ATLAS detector at the LHC*, Physics Letters B **716** (2012) 1.
- [7] CMS collaboration, S. Chatrchyan *et al.*, *Observation of a new boson at a mass of 125 GeV with the CMS experiment at the LHC*, Physics Letters B **716** (2012) 30.
- [8] F. Englert and R. Brout, *Broken Symmetry and the Mass of Gauge Vector Mesons*, Phys. Rev. Lett. **13** (1964) 321.
- [9] P. W. Higgs, *Broken symmetries, massless particles and gauge fields*, Phys. Lett. **12** (1964) 132.
- [10] P. W. Higgs, *Broken symmetries and the masses of gauge bosons*, Phys. Rev. Lett. **13** (1964) 508.
- [11] G. S. Guralnik, C. R. Hagen, and T. W. B. Kibble, *Global Conservation Laws and Massless Particles*, Phys. Rev. Lett. **13** (1964) 585.
- [12] H. Greaves and T. Thomas, *On the CPT theorem*, Studies in History and Philosophy of Science Part B: Studies in History and Philosophy of Modern Physics **45** (2014) 46.
- [13] C. S. Wu *et al.*, *Experimental Test of Parity Conservation in Beta Decay*, Phys. Rev. **105** (1957) 1413.
- [14] J. H. Christenson, J. W. Cronin, V. L. Fitch, and R. Turlay, *Evidence for the 2π Decay of the K_2^0 Meson*, Phys. Rev. Lett. **13** (1964) 138.
- [15] E. Noether, *Invariante Variationsprobleme*, Nachrichten von der Gesellschaft der Wissenschaften zu Göttingen, Mathematisch-Physikalische Klasse **1918** (1918) 235.
- [16] D. J. Gross and F. Wilczek, *Ultraviolet Behavior of Non-Abelian Gauge Theories*, Phys. Rev. Lett. **30** (1973) 1343.
- [17] B. R. Martin and G. Shaw, *Particle Physics*, Wiley, Chichester, 2008.
- [18] LHCb collaboration, R. Aaij *et al.*, *Observation of $J/\psi p$ resonances consistent with pentaquark states in $\Lambda_b^0 \rightarrow J/\psi p K^-$ decays*, Phys. Rev. Lett. **115** (2015) 072001, [arXiv:1507.03414](#).

Bibliography

- [19] L.-L. Chau and W.-Y. Keung, *Comments on the Parametrization of the Kobayashi-Maskawa Matrix*, Phys. Rev. Lett. **53** (1984) 1802.
- [20] L. Wolfenstein, *Parametrization of the Kobayashi-Maskawa Matrix*, Phys. Rev. Lett. **51** (1983) 1945.
- [21] C. Jarlskog, *Commutator of the Quark Mass Matrices in the Standard Electroweak Model and a Measure of Maximal CP Nonconservation*, Phys. Rev. Lett. **55** (1985) 1039.
- [22] L. Vale Silva, *2023 update of the extraction of the CKM matrix elements*, in *12th International Workshop on the CKM Unitarity Triangle*, 2024, arXiv:2405.08046.
- [23] J. Charles *et al.*, *CP violation and the CKM matrix: assessing the impact of the asymmetric B factories*, The Eur. Phys. J. C **41** (2005) 1.
- [24] S. L. Glashow, J. Iliopoulos, and L. Maiani, *Weak Interactions with Lepton-Hadron Symmetry*, Phys. Rev. D **2** (1970) 1285.
- [25] ATLAS collaboration, *Observation of a cross-section enhancement near the $t\bar{t}$ production threshold in $\sqrt{s} = 13$ TeV pp collisions with the ATLAS detector*, ATLAS-CONF-2025-008, 2025.
- [26] CMS collaboration, A. Hayrapetyan *et al.*, *Observation of a pseudoscalar excess at the top quark pair production threshold*, Rept. Prog. Phys. **88** (2025) 087801, arXiv:2503.22382.
- [27] A. Lenz, *Theory of Mixing and CP violation*, PoS **LHCP2018** (2018) 174, arXiv:1809.09452.
- [28] J. Albrecht, F. Bernlochner, A. Lenz, and A. Rusov, *Lifetimes of b-hadrons and mixing of neutral B-mesons: theoretical and experimental status*, Eur. Phys. J. ST **233** (2024) 359, arXiv:2402.04224.
- [29] LHCb collaboration, R. Aaij *et al.*, *Precise determination of the B_s^0 - \bar{B}_s^0 oscillation frequency*, Nature Physics **18** (2022) 1, arXiv:2104.04421.
- [30] LHCb collaboration, R. Aaij *et al.*, *Evidence for CP violation in time-integrated $D^0 \rightarrow h^- h^+$ decay rates*, Phys. Rev. Lett. **108** (2012) 111602, arXiv:1112.0938.
- [31] LHCb collaboration, R. Aaij *et al.*, *Measurement of the semileptonic CP asymmetry in B^0 - \bar{B}^0 mixing*, Phys. Rev. Lett. **114** (2015) 041601, arXiv:1409.8586.
- [32] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP violation in $B^0 \rightarrow \psi(\rightarrow \ell^+ \ell^-) K_S^0(\rightarrow \pi^+ \pi^-)$ decays*, arXiv:2309.09728, submitted to Phys. Rev. Lett.
- [33] M. Gronau, *Direct CP violation in B decays*, in *5th Conference on Flavor Physics and CP Violation*, **C070512**, 007, 2007, arXiv:0706.2156.
- [34] N. Tuning, *Lecture notes on CP violation*, 2020. <https://www.nikhef.nl/h71/Lectures/2015/ppII-cpviolation-29012015.pdf>.
- [35] J. C. Collins, D. E. Soper, and G. F. Sterman, *Factorization of Hard Processes in QCD*, Adv. Ser. Direct. High Energy Phys. **5** (1989) 1, arXiv:hep-ph/0409313.
- [36] M. Beneke, G. Buchalla, M. Neubert, and C. T. Sachrajda, *QCD factorization for exclusive non-leptonic B-meson decays: general arguments and the case of heavy-light final states*, Nuclear Physics B **591** (2000) 313.

- [37] T. Gershon, A. Lenz, A. V. Rusov, and N. Skidmore, *Testing the Standard Model with CP asymmetries in flavor-specific nonleptonic decays*, Physical Review D **105** (2022) .
- [38] M. Bordone *et al.*, *A puzzle in $\bar{B}_{(s)}^0 \rightarrow D_{(s)}^{(*)+} \{\pi^-, K^-\}$ decays and extraction of the f_s/f_d fragmentation fraction*, The Eur. Phys. J. C **80** (2020) .
- [39] S. Descotes-Genon, L. Hofer, J. Matias, and J. Virto, *Global analysis of $b \rightarrow s\ell\ell$ anomalies*, JHEP **2016** (2016) 92.
- [40] L.-S. Geng *et al.*, *Implications of new evidence for lepton-universality violation in $b \rightarrow s\ell^+\ell^-$ decays*, Phys. Rev. D **104** (2021) 035029.
- [41] S. Descotes-Genon, D. Ghosh, J. Matias, and M. Ramon, *Exploring new physics in the C \mathcal{T} - C \mathcal{T} plane*, JHEP **2011** (2011) .
- [42] G. Hiller and I. Nišandžić, *R_K and R_K^* beyond the standard model*, Physical Review D **96** (2017) .
- [43] N. Aghanim *et al.*, *Planck2018 results: I. Overview and the cosmological legacy of Planck*, Astronomy & Astrophysics **641** (2020) A1.
- [44] A. D. Sakharov, *Violation of CP Invariance, C asymmetry, and baryon asymmetry of the universe*, Pisma Zh. Eksp. Teor. Fiz. **5** (1967) 32.
- [45] G. Alonso-Álvarez, G. Elor, and M. Escudero, *Collider signals of baryogenesis and dark matter from B mesons: A roadmap to discovery*, Phys. Rev. D **104** (2021) 035028.
- [46] SNO collaboration, Q. R. Ahmad *et al.*, *Measurement of the Rate of $\nu_e + d \rightarrow p + p + e^-$ Interactions Produced by 8B Solar Neutrinos at the Sudbury Neutrino Observatory*, Phys. Rev. Lett. **87** (2001) 071301.
- [47] K. M. Case, *Reformulation of the Majorana Theory of the Neutrino*, Phys. Rev. **107** (1957) 307.
- [48] S. Gariazzo *et al.*, *Neutrino mass and mass ordering: no conclusive evidence for normal ordering*, Journal of Cosmology and Astroparticle Physics **2022** (2022) 010.
- [49] M. Chaichian, R. G. Felipe, and K. Huitu, *On quadratic divergences and the Higgs mass*, Physics Letters B **363** (1995) 101.
- [50] G. Altarelli, *The Higgs: so simple yet so unnatural*, Phys. Scripta T **158** (2013) 014011, [arXiv:1308.0545](https://arxiv.org/abs/1308.0545).
- [51] S. Iguro and T. Kitahara, *Implications for new physics from a novel puzzle in $\bar{B}_{(s)}^0 \rightarrow D_{(s)}^{(*)+} \pi^-, K^-$ decays*, Phys. Rev. D **102** (2020) 071701.
- [52] P. Koppenburg, *Patrick Koppenburg—Flavour (non-)Anomalies*, <https://www.koppenburg.ch/anomalies.html>. Accessed: 03.06.2025.
- [53] Heavy Flavor Averaging Group (HFLAV), S. Banerjee *et al.*, *Averages of b-hadron, c-hadron, and τ -lepton properties as of 2023*, Phys. Rev. D **113** (2026) 012008, [arXiv:2411.18639](https://arxiv.org/abs/2411.18639).

Bibliography

- [54] E. Lopienska, *The CERN accelerator complex, layout in 2022. Complexe des accélérateurs du CERN en janvier 2022*, 2022. General Photo, <https://cds.cern.ch/record/2800984>.
- [55] M. Vretenar *et al.*, *Linac4 design report*, vol. 6 of *CERN Yellow Reports: Monographs*, CERN, Geneva, 2020.
- [56] M. Benedikt *et al.*, *LHC Design Report*, CERN Yellow Reports: Monographs, CERN, Geneva, 2004.
- [57] LHCb collaboration, R. Aaij *et al.*, *Measurement of the b -quark production cross-section in 7 and 13 TeV pp collisions*, Phys. Rev. Lett. **118** (2017) 052002, Erratum *ibid.* **119** (2017) 169901, arXiv:1612.05140.
- [58] LHCb collaboration, C. Elsasser, *$\bar{b}b$ production angle plots*, https://lhcb.web.cern.ch/lhcb/speakersbureau/html/bb_ProductionAngles.html. Accessed: 30. August 2025.
- [59] R. Alemany-Fernandez, F. Follin, and R. Jacobsson, *The LHCb Online Luminosity Control and Monitoring*, in *4th International Particle Accelerator Conference*, TUPFI010, 2013.
- [60] C. O. Dib, C. S. Kim, and N. A. Neill, *The b quark fragmentation fractions at LHCb and meson decays with heavy quark spectators*, The Eur. Phys. J. C **83** (2023) 793.
- [61] LHCb collaboration, *LHCb Tracker Upgrade Technical Design Report*, CERN-LHCC-2014-001, 2014.
- [62] 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.
- [63] LHCb collaboration, *LHCb magnet: Technical Design Report*, CERN-LHCC-2000-007, 2000.
- [64] LHCb collaboration, *LHCb VELO (Vertex Locator): Technical Design Report*, CERN-LHCC-2001-011, 2001.
- [65] R. Aaij *et al.*, *Performance of the LHCb Vertex Locator*, JINST **9** (2014) P09007, arXiv:1405.7808.
- [66] LHCb collaboration, A. C. S. Davis, *Performance of the LHCb tracking system in Run I of the LHC*, Nucl. Instrum. Methods Phys. Res. , A **824** (2016) 15.
- [67] LHCb collaboration, F. Dordei, *LHCb detector and trigger performance in Run II*, EPJ Web Conf. **164** (2017) 01016.
- [68] S. Haywood, *Impact Parameter Resolution in the Presence of Multiple- Scattering*, ATL-INDET-94-091, Geneva, 1994.
- [69] C. Ilgner *et al.*, *The Beam Conditions Monitor of the LHCb Experiment*, arXiv:1001.2487.
- [70] LHCb collaboration, *LHCb online system, data acquisition and experiment control: Technical Design Report*, CERN-LHCC-2001-040, 2001.
- [71] LHCb collaboration, *LHCb inner tracker: Technical Design Report*, CERN-LHCC-2002-029, 2002.

- [72] R. Arink *et al.*, *Performance of the LHCb Outer Tracker*, JINST **9** (2014) P01002, arXiv:1311.3893.
- [73] LHCb collaboration, *LHCb RICH: Technical Design Report*, CERN-LHCC-2000-037, 2000.
- [74] M. Adinolfi *et al.*, *Performance of the LHCb RICH detector at the LHC*, Eur. Phys. J. **C73** (2013) 2431, arXiv:1211.6759.
- [75] LHCb collaboration, *LHCb calorimeters: Technical Design Report*, CERN-LHCC-2000-036, 2000.
- [76] 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.
- [77] LHCb collaboration, R. Aaij *et al.*, *Measurement of lepton universality parameters in $B^+ \rightarrow K^+ \ell^+ \ell^-$ and $B^0 \rightarrow K^{*0} \ell^+ \ell^-$ decays*, Phys. Rev. **D108** (2023) 032002, arXiv:2212.09153.
- [78] LHCb collaboration, R. Aaij *et al.*, *Test of lepton universality in $b \rightarrow s \ell^+ \ell^-$ decays*, Phys. Rev. Lett. **131** (2023) 051803, arXiv:2212.09152.
- [79] LHCb collaboration, *LHCb muon system: Technical Design Report*, CERN-LHCC-2001-010, 2001.
- [80] D. Derkach *et al.*, *Machine-Learning-based global particle-identification algorithms at the LHCb experiment*, Journal of Physics: Conference Series **1085** (2018) 042038.
- [81] LHCb collaboration, *Global PID performance for charged particles*, LHCb-FIGURE-2020-012, 2020.
- [82] J. Albrecht *et al.*, *Summary of the trigger systems of the Large Hadron Collider experiments ALICE, ATLAS, CMS and LHCb*, Journal of Physics G: Nuclear and Particle Physics **52** (2025) 030501.
- [83] C. A. Aidala *et al.*, *Ntuple Wizard: An Application to Access Large-Scale Open Data from LHCb*, Comput. Softw. Big Sci. **7** (2023) 6.
- [84] LHCb collaboration, J. Albrecht, *The LHCb Trigger System*, Nucl. Phys. B, Proc. Suppl. **187** (2009) 237.
- [85] P. Astier *et al.*, *Kalman filter track fits and track breakpoint analysis*, Nucl. Instrum. Meth. **450** (2000) 138.
- [86] S. Benson, V. Gligorov, M. A. Vesterinen, and J. M. Williams, *The LHCb Turbo Stream*, J. Phys. : Conf. Ser. **664** (2015) 082004.
- [87] R. Aaij *et al.*, *A comprehensive real-time analysis model at the LHCb experiment*, JINST **14** (2019) P04006, arXiv:1903.01360.
- [88] LHCb collaboration, *LHCb computing: Technical Design Report*, CERN-LHCC-2005-019, 2005.
- [89] R. Brun and F. Rademakers, *ROOT — An object oriented data analysis framework*, Nucl. Instrum. Meth. **389** (1997) 81, New Computing Techniques in Physics Research V.

Bibliography

- [90] N. A. Grieser *et al.*, *The LHCb Stripping Project: Sustainable Legacy Data Processing for High-Energy Physics*, *Comput. Softw. Big Sci.* **9** (2025) 21, [arXiv:2509.05294](#).
- [91] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $B_s^0 \rightarrow D_s^\mp K^\pm$ decays*, *JHEP* **2503** (2025) 139, [arXiv:2412.14074](#).
- [92] LHCb collaboration, D. Fazzini, *Flavour Tagging in the LHCb experiment*, *PoS LHCP2018* (2018) 230.
- [93] S. Miglioranza *et al.*, *The LHCb Simulation Application, Gauss: Design, Evolution and Experience*, CERN, Geneva, 2011.
- [94] T. Sjöstrand, S. Mrenna, and P. Skands, *A brief introduction to PYTHIA 8.1*, *Computer Physics Communications* **178** (2008) 852.
- [95] T. Sjöstrand, S. Mrenna, and P. Skands, *PYTHIA 6.4 physics and manual*, *JHEP* **2006** (2006) 026.
- [96] LHCb collaboration, I. Belyaev *et al.*, *Handling of the generation of primary events in Gauss, the LHCb simulation framework*, in *18th International Conference on Computing in High Energy and Nuclear Physics* (S. C. Lin, ed.), **331** 032047, 2011.
- [97] D. J. Lange, *The EvtGen particle decay simulation package*, *Nucl. Instrum. Meth. A* **462** (2001) 152.
- [98] N. Davidson, T. Przedzinski, and Z. Was, *PHOTOS Interface in C++; Technical and Physics Documentation*, 2015.
- [99] GEANT4 collaboration, S. Agostinelli *et al.*, *GEANT4—a simulation toolkit*, *Nucl. Instrum. Methods Phys. Res., A* **506** (2003) 250.
- [100] V. V. Gligorov, *Reconstruction of the Channel $B_d^0 \rightarrow D^+ \pi^-$ and Background Classification at LHCb (revised)*, LHCb-2007-044, Geneva, 2007.
- [101] LHCb collaboration, R. Aaij *et al.*, *Updated measurement of time-dependent CP-violating observables in $B_s^0 \rightarrow J/\psi K^+ K^-$ decays*, *Eur. Phys. J.* **C79** (2019) 706, Erratum *ibid.* **C80** (2020) 601, [arXiv:1906.08356](#).
- [102] M. Artuso, G. Borissov, and A. Lenz, *CP violation in the B_s^0 system*, *Rev. Mod. Phys.* **88** (2016) 045002.
- [103] R. H. Dalitz, *CXII. On the analysis of τ -meson data and the nature of the τ -meson*, *The London, Edinburgh, and Dublin Philosophical Magazine and Journal of Science* **44** (1953) 1068, [arXiv:https://doi.org/10.1080/14786441008520365](#).
- [104] L. Anderlini *et al.*, *The PIDCalib package*, LHCb-PUB-2016-021, Geneva, 2016.
- [105] M. Pivk and F. R. Le Diberder, *sPlot: A statistical tool to unfold data distributions*, *Nucl. Instrum. Meth.* **555** (2005) 356.
- [106] A. Martin Sanchez, P. Robbe, and M.-H. Schune, *Performances of the LHCb L0 Calorimeter Trigger*, LHCb-PUB-2011-026, Geneva, 2012.
- [107] S. Monteil, P. Robbe, and T. Miralles, *Determination of the L0Hadron trigger efficiency corrections*, LHCb-INT-2024-013, Geneva, 2024.
- [108] LHCb collaboration, R. Aaij *et al.*, *Measurement of $\sigma(pp \rightarrow b\bar{b}X)$ at $\sqrt{s} = 7$ TeV in the forward region*, *Phys. Lett.* **B694** (2010) 209, [arXiv:1009.2731](#).

- [109] LHCb collaboration, R. Aaij *et al.*, *Measurement of the track reconstruction efficiency at LHCb*, JINST **10** (2015) P02007, arXiv:1408.1251.
- [110] K. Cranmer, *Practical Statistics for the LHC*, in *2011 European School of High-Energy Physics*, 267–308, 2014, arXiv:1503.07622.
- [111] J. Amoraal *et al.*, *Application of vertex and mass constraints in track-based alignment*, Nucl. Instrum. Meth. A **712** (2013) 48, arXiv:1207.4756.
- [112] LHCb collaboration, R. Aaij *et al.*, *Precise measurement of the f_s/f_d ratio of fragmentation fractions and of B_s^0 decay branching fractions*, Phys. Rev. **D104** (2021) 032005, arXiv:2103.06810.
- [113] R. J. Barlow, *Practical statistics for particle physics*, CERN Yellow Rep. School Proc. **5** (2020) 149, arXiv:1905.12362.
- [114] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP observables in $B^\pm \rightarrow D^{(*)}K^\pm$ and $B^\pm \rightarrow D^{(*)}\pi^\pm$ decays using two-body D final states*, JHEP **04** (2021) 081, arXiv:2012.09903.
- [115] LHCb collaboration, R. Aaij *et al.*, *Measurement of the CP asymmetry in $B_s^0-\bar{B}_s^0$ mixing*, Phys. Rev. Lett. **117** (2016) 061803, arXiv:1605.09768.
- [116] J. Davies, *Beauty to Open Charm Measurements at LHCb and Other Flavourful Friends*, PhD thesis, 2024, Manchester U., 2024.
- [117] A. Davis *et al.*, *Measurement of the instrumental asymmetry for $K^-\pi^+$ -pairs at LHCb in Run 2*, LHCb-PUB-2018-004, Geneva, 2018.
- [118] LHCb collaboration, R. Aaij *et al.*, *Measurement of CP asymmetry in $D^0 \rightarrow K^-K^+$ and $D^0 \rightarrow \pi^-\pi^+$ decays*, JHEP **07** (2014) 041, arXiv:1405.2797.
- [119] A. Rogozhnikov, *Reweighting with Boosted Decision Trees*, Journal of Physics: Conference Series **762** (2016) 012036.
- [120] L. Dufour, *High-precision measurements of charge asymmetries at LHCb*, PhD thesis, Groningen U., 2019.
- [121] LHCb collaboration, *Framework TDR for the LHCb Upgrade: Technical Design Report*, CERN-LHCC-2012-007, 2012.
- [122] LHCb collaboration, R. Aaij *et al.*, *The LHCb Upgrade I*, JINST **19** (2024) P05065, arXiv:2305.10515.
- [123] LHCb collaboration, *LHCb VELO Upgrade Technical Design Report*, CERN-LHCC-2013-021, 2013.
- [124] LHCb collaboration, *Primary Vertex resolution with early 2024 data*, LHCb-FIGURE-2024-011, 2024.
- [125] LHCb collaboration, *Track Momentum Resolution at LHCb in 2024*, LHCb-FIGURE-2024-040, 2024.
- [126] LHCb RICH collaboration, A. Papanestis, F. Keizer, and S. A. Wotton, *The upgrade of the LHCb RICH system for the LHC Run 3*, JINST **15** (2020) C09022.
- [127] LHCb collaboration, *Hadron PID Performance in 2024*, LHCb-FIGURE-2024-031, 2024.

Bibliography

- [128] LHCb collaboration, *LHCb PLUME: Probe for LUMinosity MEasurementt*, CERN-LHCC-2021-002, 2021.
- [129] LHCb collaboration, R. Aaij *et al.*, *A Comparison of CPU and GPU implementations for the LHCb Experiment Run 3 Trigger*, *Comput. Softw. Big Sci.* **6** (2021) 1, [arXiv:2105.04031](#), 30 pages, 15 figures, 8 tables.
- [130] LHCb collaboration, *RTA and DPA dataflow diagrams for Run 1, Run 2, and the upgraded LHCb detector*, LHCb-FIGURE-2020-016, 2020.
- [131] K. Wyllie *et al.*, *Electronics Architecture of the LHCb Upgrade*, CERN, Geneva, 2013.
- [132] J. Peddie, *The History of the GPU - New Developments*, Springer International Publishing, Cham, 2022.
- [133] R. Aaij *et al.*, *Allen: A High-Level Trigger on GPUs for LHCb*, *Comput. Softw. Big Sci.* **4** (2020) .
- [134] M. Herlihy and N. Shavit, *The Art of Multiprocessor Programming, Revised Reprint*, Morgan Kaufmann Publishers Inc., San Francisco, CA, USA, 1st ed., 2012.
- [135] A. Bailly-Reyre *et al.*, *Looking Forward: A High-Throughput Track Following Algorithm for Parallel Architectures*, *IEEE Access* **12** (2024) 114198, [arXiv:2402.14670](#).
- [136] T. Evans, C. Fitzpatrick, and J. Horswill, *An automated bandwidth division for the LHCb upgrade trigger*, *Comput. Softw. Big Sci.* **9** (2025) 7, [arXiv:2502.09557](#).
- [137] V. Kholoimov *et al.*, *A Downstream and vertexing algorithm for Long Lived Particles (LLP) selection at the first High-level trigger (HLT1) of LHCb*, 2025.
- [138] Valls Canudas, Núria, Vilasis Cardona, Xavier, Calvo Gómez, Míriam, and Golobardes Ribé, Elisabet, *Preliminary Performance Study of an Alternative Calorimeter Clustering Solution for Allen in LHCb*, *EPJ Web of Conf.* **295** (2024) 02023.
- [139] F. Reiss, *Real-time alignment procedure at the LHCb experiment for Run 3*, in *7th International Connecting The Dots Workshop*, 2023.
- [140] M. Adinolfi *et al.*, *LHCb data quality monitoring*, *J. Phys. : Conf. Ser.* **898** (2017) 092027.
- [141] C. Bozzi, *LHCb Computing Resources: preliminary 2025 requests*, LHCb-PUB-2023-003, Geneva, 2023.
- [142] N. Skidmore, E. Rodrigues, and P. Koppenburg, *Run-3 offline data processing and analysis at LHCb*, *PoS EPS-HEP2021* (2022) 792.
- [143] A. Abdelmotteleb *et al.*, *The LHCb Sprucing and Analysis Productions*, *Comput. Softw. Big Sci.* **9** (2025) 15.
- [144] A. Mathad *et al.*, *FunTuple: A New N-tuple Component for Offline Data Processing at the LHCb Experiment*, *Comput. Softw. Big Sci.* **8** (2024) 6.
- [145] Burr, Chris, Couturier, Ben, and O’Neil, Ryunosuke, *Facilitating the preservation of LHCb Analyses with APD*, *EPJ Web of Conf.* **295** (2024) 08008.

- [146] M. Mazurek, M. Clemencic, and G. Corti, *Gauss and Gaussino: the LHCb simulation software and its new experiment agnostic core framework*, PoS **ICHEP2022** (2022) 225.
- [147] M. Frank, F. Gaede, C. Grefe, and P. Mato, *DD4hep: A Detector Description Toolkit for High Energy Physics Experiments*, Journal of Physics: Conference Series **513** (2014) 022010.
- [148] LHCb collaboration, R. Aaij *et al.*, *Measurements of charmed meson and antimeson production asymmetries at $\sqrt{s} = 13.6$ TeV*, JHEP **10** (2025) 050, [arXiv:2505.14494](https://arxiv.org/abs/2505.14494).
- [149] D. Friday, *The LHCb VELO detector: Design, operation and first results*, Nucl. Instrum. Meth. **1070** (2025) 170028.
- [150] LHCb collaboration, LHCb collaboration, *LHCb Operations Plots Web-page*, <https://lbggroups.cern.ch/online/OperationsPlots/index.htm> Accessed: 30.08.2025.
- [151] P. Jaccard, *The Distribution of the Flora in the Alpine Zone*, New Phytologist **11** (1912) 37.
- [152] L. Grazette *et al.*, *A Comprehensive Bandwidth Testing Framework for the LHCb Upgrade Trigger System*, in *27th International Conference on Computing in High Energy and Nuclear Physics*, **337** 01261, 2025, [arXiv:2503.19582](https://arxiv.org/abs/2503.19582).
- [153] Hunter, Ross *et al.*, *An automated tool to facilitate consistent test-driven development of trigger selections for LHCb's Run 3*, EPJ Web Conf. **251** (2021) 04024.
- [154] O. Kitouni, N. Nolte, and M. Williams, *Robust and provably monotonic networks*, Machine Learning: Science and Technology **4** (2023) 035020.
- [155] LHCb collaboration, R. Aaij *et al.*, *Observation of the very rare $\Sigma^+ \rightarrow p\mu^+\mu^-$ decay*, Phys. Rev. Lett. **135** (2025) 051801, [arXiv:2504.06096](https://arxiv.org/abs/2504.06096), Submitted to Phys. Rev. Lett.
- [156] M. D. Schwartz, *Modern Machine Learning and Particle Physics*, Harvard Data Science Review (2021) .
- [157] V. V. Gligorov and M. Williams, *Efficient, reliable and fast high-level triggering using a bonsai boosted decision tree*, JINST **8** (2013) P02013.
- [158] N. Schulte *et al.*, *Development of the Topological Trigger for LHCb Run 3*, 2023.
- [159] A. Mao, M. Mohri, and Y. Zhong, *Cross-Entropy Loss Functions: Theoretical Analysis and Applications*, 2023.
- [160] R. T. Marler and J. S. Arora, *The weighted sum method for multi-objective optimization: new insights*, Structural and Multidisciplinary Optimization **41** (2010) 853.
- [161] G. Kasieczka and D. Shih, *Robust Jet Classifiers through Distance Correlation*, Physical Review Letters **125** (2020) .
- [162] H. Dembinski *et al.*, *scikit-hep/iminuit*, 2025. doi: 10.5281/zenodo.15157028.
- [163] M. Hatlo *et al.*, *Developments of mathematical software libraries for the LHC experiments*, IEEE Trans. Nucl. Sci. **52** (2005) 2818.

Bibliography

- [164] J. Albrecht *et al.*, *Triggercalib: a turnkey package for estimating lhcb trigger efficiencies*, [arXiv:2505.15951](#).
- [165] S. Tolk, J. Albrecht, F. Dettori, and A. Pellegrino, *Data driven trigger efficiency determination at LHCb*, CERN, Geneva, 2014.
- [166] W. Verkerke and D. Kirkby, in *The RooFit Toolkit for Data Modeling*, pp. 186–189.
- [167] J. Eschle, A. Puig Navarro, R. Silva Coutinho, and N. Serra, *zfit: Scalable pythonic fitting*, *SoftwareX* **11** (2020) 100508.
- [168] H. Dembinski, M. Kenzie, C. Langenbruch, and M. Schmelling, *Custom Orthogonal Weight functions (COWs) for event classification*, *Nucl. Instrum. Meth.* **1040** (2022) 167270.
- [169] M. G. Kendall, *Rank Correlation Methods*, Theory and applications of rank order-statistics, Griffin, 1970.
- [170] M. Paterno, *Calculating efficiencies and their uncertainties*, doi: 10.2172/15017262.
- [171] E. B. W. and, *Probable Inference, the Law of Succession, and Statistical Inference*, *Journal of the American Statistical Association* **22** (1927) 209.
- [172] H. Dembinski and M. Schmelling, *Bias, variance, and confidence intervals for efficiency estimators in particle physics experiments*, 2022.
- [173] LHCb collaboration, *HLT1 trigger efficiencies in 2024 data*, LHCb-FIGURE-2024-030, 2024.
- [174] LHCb collaboration, *Inclusive trigger efficiencies in 2025 data*, LHCb-FIGURE-2025-015, 2025.
- [175] LHCb collaboration, *LHCb Trigger and Online Upgrade Technical Design Report*, CERN-LHCC-2014-016, 2014.
- [176] E. Rodrigues, *Dealing with clones in the tracking*, CERN, Geneva, 2006.
- [177] K. Cranmer, *Kernel estimation in high-energy physics*, *Computer Physics Communications* **136** (2001) 198.
- [178] B. Efron and R. J. Tibshirani, *An introduction to the bootstrap*, Chapman & Hall/CRC monographs on statistics and applied probability, Chapman and Hall, London, 1993.
- [179] LHCb collaboration, R. Aaij *et al.*, *First measurement of the decay-time-integrated CP asymmetry in $B_s^0 \rightarrow D_s^- \pi^+$ decays*, [arXiv:2603.10860](#).
- [180] M. Oreglia, *A Study of the Reactions $\psi' \rightarrow \gamma\gamma\psi$* , PhD thesis, 1980.
- [181] D. Martinez Santos and F. Dupertuis, *Mass distributions marginalized over per-event errors*, *Nucl. Instrum. Meth.* **764** (2014) 150.
- [182] LHCb collaboration, R. Aaij *et al.*, *Measurement of the CKM angle γ in $B^\pm \rightarrow DK^\pm$ and $B^\pm \rightarrow D\pi^\pm$ decays with $D \rightarrow K_S^0 h^+ h^-$* , *JHEP* **02** (2021) 0169, [arXiv:2010.08483](#).
- [183] N. L. Johnson, *Systems of Frequency Curves Generated by Methods of Translation*, *Biometrika* **36** (1949) 149.

- [184] LHCb collaboration, *LHCb RunDB*, <https://lbrundb.cern.ch> Accessed: 03.09.2025.